Energy Harvesting Networked Nodes: Measurements, Algorithms, and Prototyping

Maria Gorlatova

Submitted in partial fulfillment of the requirements for the degree of Doctor of Philosophy in the Graduate School of Arts and Sciences.

COLUMBIA UNIVERSITY

2013

© 2013 Maria Gorlatova All rights reserved

ABSTRACT

Energy Harvesting Networked Nodes: Measurements, Algorithms, and Prototyping

Maria Gorlatova

Recent advances in ultra-low-power wireless communications and in energy harvesting will soon enable *energetically self-sustainable wireless devices*. Networks of such devices will serve as building blocks for different *Internet of Things* (IoT) applications, such as searching for an object on a network of objects and continuous monitoring of object configurations. Yet, numerous challenges need to be addressed for the IoT vision to be fully realized.

This thesis considers several challenges related to ultra-low-power energy harvesting networked nodes: *energy source characterization*, *algorithm design*, and *node design and prototyping*. Additionally, the thesis contributes to engineering education, specifically to *project-based learning*.

We summarize our contributions to *light and kinetic (motion) energy characterization* for energy harvesting nodes. To characterize light energy, we conducted a first-of-its kind 16 month-long indoor light energy measurements campaign. To characterize energy of motion, we collected over 200 hours of human and object motion traces. We also analyzed traces previously collected in a study with over 40 participants. We summarize our insights, including light and motion energy budgets, variability, and influencing factors. These insights are useful for designing energy harvesting nodes and energy harvesting adaptive algorithms. We shared with the community our light energy traces, which can be used as energy inputs to system and algorithm simulators and emulators.

We also discuss *resource allocation problems* we considered for energy harvesting nodes. Inspired by the needs of tracking and monitoring IoT applications, we formulated and studied resource allocation problems aimed at allocating the nodes' time-varying resources in a uniform way with respect to time. We mainly considered deterministic energy profile and stochastic environmental energy models, and focused on single node and link scenarios. We formulated optimization problems using utility maximization and lexicographic maximization frameworks, and introduced algorithms for solving the formulated problems. For several settings, we provided low-complexity solution algorithms. We also examined many simple policies. We demonstrated, analytically and via simulations, that in many settings simple policies perform well.

We also summarize our design and prototyping efforts for a new class of ultra-low-power nodes – Energy Harvesting Active Networked Tags (EnHANTs). Future EnHANTs will be wireless nodes that can be attached to commonplace objects (books, furniture, clothing). We describe the En-HANTs prototypes and the EnHANTs testbed that we developed, in collaboration with other research groups, over the last 4 years in 6 integration phases. The prototypes harvest energy of the indoor light, communicate with each other via ultra-low-power transceivers, form small multihop networks, and adapt their communications and networking to their energy harvesting states. The EnHANTs testbed can expose the prototypes to light conditions based on real-world light energy traces. Using the testbed and our light energy traces, we evaluated some of our energy harvesting adaptive policies. Our insights into node design and performance evaluations may apply beyond EnHANTs to networks of various energy harvesting nodes.

Finally, we present our contributions to *engineering education*. Over the last 4 years, we engaged high school, undergraduate, and M.S. students in more than 100 research projects within the EnHANTs project. We summarize our approaches to facilitating student learning, and discuss the results of evaluation surveys that demonstrate the effectiveness of our approaches.

Table of Contents

1	Intr	roducti	on	1
	1.1	Energy	$^{\prime}$ Harvesting and Storage for Ultra-low-power Energy Harvesting Nodes	3
	1.2	Design	Space for Energy Harvesting Adaptive Algorithms	4
	1.3	Energy	Harvesting Active Networked Tags (EnHANTs)	7
	1.4	Summ	ary of Contributions	9
		1.4.1	Characterizing Environmental Energy for Energy Harvesting Nodes $\ \ . \ . \ .$	10
		1.4.2	Resource Allocation Algorithms for Energy Harvesting Nodes	11
		1.4.3	EnHANTs Prototypes Testbed for Energy Harvesting Adaptive Policy Evalu-	
			ations	12
		1.4.4	Project-based Learning	12
2	Rel	ated W	⁷ ork	14
2	Rel 2.1	ated W Energy	York 7 Source Characterization for Ultra-low-power Energy Harvesting Nodes	14 14
2	Rel 2.1 2.2	ated W Energy Energy	Vork 7 Source Characterization for Ultra-low-power Energy Harvesting Nodes 7 Harvesting Adaptive Algorithms	14 14 15
2	Rel 2.1 2.2 2.3	ated W Energy Energy Design	Vork 7 Source Characterization for Ultra-low-power Energy Harvesting Nodes 7 Harvesting Adaptive Algorithms	 14 14 15 17
2	Rel 2.1 2.2 2.3 2.4	ated W Energy Energy Design Projec	Vork v Source Characterization for Ultra-low-power Energy Harvesting Nodes v Harvesting Adaptive Algorithms ing and Prototyping Energy Harvesting Active Networked Tags t-based Learning	 14 14 15 17 18
2 3	 Rel 2.1 2.2 2.3 2.4 Cha 	ated W Energy Energy Design Projec	Vork v Source Characterization for Ultra-low-power Energy Harvesting Nodes v Harvesting Adaptive Algorithms ing and Prototyping Energy Harvesting Active Networked Tags t-based Learning t-based Learning traiting Light Energy for Energy Harvesting Nodes	 14 14 15 17 18 20
2 3	Rel 2.1 2.2 2.3 2.4 Cha 3.1	ated W Energy Design Projec aracter Metho	Vork v Source Characterization for Ultra-low-power Energy Harvesting Nodes v Harvesting Adaptive Algorithms ing and Prototyping Energy Harvesting Active Networked Tags t-based Learning tizing Light Energy for Energy Harvesting Nodes dology	 14 14 15 17 18 20 21
2	Rel 2.1 2.2 2.3 2.4 Cha 3.1 3.2	ated W Energy Design Projec aracter Metho Spatia	Vork v Source Characterization for Ultra-low-power Energy Harvesting Nodes v Harvesting Adaptive Algorithms ing and Prototyping Energy Harvesting Active Networked Tags t-based Learning tizing Light Energy for Energy Harvesting Nodes dology I Variability	 14 14 15 17 18 20 21 22
3	Rel 2.1 2.2 2.3 2.4 Cha 3.1 3.2 3.3	ated W Energy Design Projec aracter Metho Spatia Node J	Vork v Source Characterization for Ultra-low-power Energy Harvesting Nodes v Harvesting Adaptive Algorithms ing and Prototyping Energy Harvesting Active Networked Tags t-based Learning t-based Learning izing Light Energy for Energy Harvesting Nodes dology I Variability Energy Budgets and Daily Energy Availability	 14 14 15 17 18 20 21 22 24

	3.5	Mobile	e Measurements	28
	3.6	Conclu	usions	29
4	Cha	aracter	izing Kinetic Energy for Energy Harvesting Nodes	30
	4.1	Model	and Measurement Setup	31
		4.1.1	Inertial Harvester Model	31
		4.1.2	Collecting Motion Information	33
		4.1.3	Harvesting Rates and Data Rates	34
		4.1.4	Optimizing Harvester Parameters	35
	4.2	Object	t Motion Energy	36
	4.3	Huma	n Motion Energy	37
		4.3.1	Study Summary	38
		4.3.2	Energy for Different Activities	39
		4.3.3	Consistency of Dominant Motion Frequency	40
		4.3.4	Dependency on Human Height and Weight $\ldots \ldots \ldots \ldots \ldots \ldots \ldots$	41
	4.4	Long-t	term Human Mobility	42
		4.4.1	Prolonged Activities	43
		4.4.2	Day-Long Human Mobility	44
		4.4.3	Harvesting Process Variability and Properties	45
	4.5	Conclu	usions and Future Work	47
5	Res	ource	Allocation Algorithms for Energy Harvesting Nodes	48
	5.1	Model		49
	5.2	Deterr	ninistic Energy Profile	52
		5.2.1	Single Node: Optimizing Node Energy Spending	52
		5.2.2	Link: Optimizing Data Rates	57
		5.2.3	Numerical Results	62
	5.3	Stocha	astic Energy Models	64
		5.3.1	Optimal Policies and Discretization Bounds	64
		5.3.2	Solution Algorithms	66

		5.3.3	Approximate Policies and Heuristics	68
		5.3.4	Numerical Results	69
	5.4	Evalua	ating Model-free Approaches with Kinetic Energy Traces	71
	5.5	Conclu	usions and Future Work	75
	5.A	Proofs	s: Deterministic Energy Profile Environmental Energy Model	76
		5.A.1	Single Node	76
		5.A.2	Link	80
	5.B	Proofs	s: Stochastic Environmental Energy Model	82
6	EnH	IANT	Prototypes Testbed for Energy Harvesting Adaptive Policy Evalua	-
	tion	IS		83
	6.1	EnHA	NT Prototypes and Prototype Testbed	85
		6.1.1	EnHANT Prototype	85
		6.1.2	EnHANT Prototype Testbed	87
	6.2	Evalua	ating Energy Harvesting Adaptive Policies	90
		6.2.1	Evaluating Optimal and Approximate Node Policies	92
		6.2.2	Evaluating Optimal and Approximate Link Policies	93
		6.2.3	Policy Implementations for the Stochastic Energy Model	94
	6.3	Energ	y Harvesting Adaptive EnHANT Networking	95
		6.3.1	Flow Control Policies	97
		6.3.2	Collection Tree Adaptation Policies	99
	6.4	Concl	usions	100
	6.A	Phase	d Prototype and Testbed Development	101
7	Con	clusio	ns	103
\mathbf{A}	Pro	ject-ba	ased Learning within the EnHANTs Project	113
	A.1	Stude	nt Projects	114
	A.2	Projec	et Organization	116
	A.3	Lessor	ns Learned	117
	A.4	Exper	iences and Feedback	119

A.5	Conclusions																						1	12	1

List of Figures

1	A network of wireless ultra-low-power energy harvesting networked nodes	2
2	Examples of different light energy source models: deterministic profile, partially pre-	
	dictable profile, and stochastic process	5
3	An example of harvested power versus energy storage "curves" for a light energy	
	harvesting node that uses a capacitor for energy storage. \ldots	6
4	The envisioned EnHANTs	7
5	A schematic diagram of the relationships between energy parameters	21
6	A schematic diagram of the indoor irradiance measurements locations	22
7	Sample irradiance measurements	22
8	Irradiance measurements in different parts of a bookshelf	23
9	Irradiance measurements in different locations on a person	23
10	Long-term daily irradiation.	25
11	Sample energy profiles for indoor and outdoor locations	26
12	Irradiance measurements recorded by a mobile device	28
13	A second-order mass-spring system model of an inertial harvester, and the frequency	
	response magnitude for two different harvesters	31
14	Acceleration measurement unit and placements	33
15	Demonstration of obtaining the power generated by an inertial harvester from the	
	recorded acceleration.	34

16	The average power generated by an inertial harvester from the same motion for dif-	
	ferent combinations of harvester parameters	35
17	Characterization of kinetic energy for common human activities, based on a 40-	
	participant study	38
18	Motion energy characterization for a 3 hour run	43
19	Kinetic energy for normal daily human routine	44
20	Motion energy harvesting process variability: the percentage of the total energy har-	
	vested at each of the demonstrated power levels, and the percentage of time the power	
	is harvested at each of the demonstrated power levels	46
21	Node energy profile and the corresponding energy allocation vectors, obtained by	
	solving the TFLA and the TFU problems	54
22	Energy profiles of link endpoints and the corresponding data rate allocation vectors,	
	obtained by solving the LTFL and the LTFU problems	59
23	Comparison of the optimal link policy and the decoupled rate control algorithms	60
24	Energy profile and energy spending rate assignments, obtained by solving the TFLA-	
	LIN and TFU-LIN problems, and by solving the TFU problem for nonlinear energy	
	storage	62
25	Communication rates, obtained by solving the LTFL and the LTFU problems. $\ . \ .$	63
26	Optimal energy spending rates, obtained by solving the SPD problem	66
27	Optimal communication rates, obtained by solving the LSPD problem	67
28	Node scenarios with a stochastic energy model: objective function value and node	
	downtime	70
29	Examples of realizations of the SG, SL, and Scheme-LB policies: data rates and	
	battery levels	72
30	Performance of the Scheme-LB, SL, and SG policies: data rates, objective function	
	values, node ON time percentages, and energy misuse percentages	73
31	SL policy performance with traces for different participants: ON times and misused	
	energy percentages	74

32	Scheme-LB policy performance using energy traces and using the corresponding ON/OFI	F,
	Markov, and i.i.d. processes: average data rates and node ON times	74
33	Current EnHANTs prototype testbed.	84
34	A block diagram of the EnHANT prototype, and its interactions with the testbed. $% \left({{{\bf{n}}_{\rm{B}}}} \right)$.	85
35	Energy Harvesting Module: a photo and a block diagram	86
36	Solar cells integrated with EnHANT prototypes	86
37	A schematic diagram of the EnHANTs testbed.	88
38	A screenshot of the EnHANTs testbed monitoring system.	88
39	A dark box enclosure used in the EnHANTs testbed	89
40	Energy harvesting for a prototype with an a-Si solar cell and with an organic photo-	
	voltaic	89
41	Sample EnHANT prototype policy realizations: OPT, SG, and CR policies for a node	
	scenario with a deterministic profile energy model, and OPT, DRC-NOPT, DRC-SG,	
	and DRC-CR policies for a link scenario with a deterministic profile energy model	91
42	Node scenarios with a deterministic profile energy model: energy harvesting rates,	
	energy spending rates and battery levels, experimental and obtained via simulations,	
	and objective function values under the OPT and CR policies.	92
43	Link scenarios with a deterministic profile energy model: energy harvesting rates,	
	objective function values, average data rates under the DRC-CR policies, and average	
	data rates under the DRC-NOPT policies	94
44	Sample EnHANT prototype policy realizations for a node scenario with a stochastic	
	energy model: OPT, SL, and THR1 policies.	95
45	Small-scale multihop network topologies a 3-node line network and a 4-node diamond	
	network	95
46	Node energy allocation policies used as building blocks for network policies: energy	
	harvesting rates and energy spending rates under the EX, EP-1, and EP-12 policies,	
	and battery levels, experimentally measured and simulated, under the EX and EP-1	
	policies	96

47	FLEX flow control policies: energy harvesting rates, data rates under the EX and	
	EP-1 node energy allocation policies, and average data rates, under the EX policy,	
	for different values of r_c .	97
48	Topology adaptation policies: energy harvesting rates and percentage of harvested	
	energy used by the nodes	99
49	A schematic diagram of phase-based EnHANTs prototype and testbed development.	101
50	EnHANTs prototypes and testbed, Phases I-V	101
51	Student projects completed in different semesters	114
52	Students involved in the EnHANTs project.	114
53	Students semesters	115
54	Project survey results	119
55	Project survey results: EnHANTs projects compared with other activities	120

List of Tables

1	Light energy measurement setups – average daily irradiation and achievable bit rates.	24
2	Mobile light energy measurements – average irradiance and achievable bit rates	28
3	Object motion measurements.	36
4	Energy budgets and data rates based on measurements of common human activities.	39
5	Energy budgets, variability, and data rates based on collected traces for daily human	
	routines	44
6	Nomenclature.	50
7	Data rates under different collection tree adaptation policies	99

Acknowledgements

First and foremost, I would like to thank my Ph.D. advisor, professor Gil Zussman, for his guidance. Gil's foresight, energy, dedication, and unwavering high standards were the key to helping me grow as a researcher and a professional. I am extraordinarily lucky to have had the chance to work with Gil, and will be forever grateful for his support and mentorship.

During my Ph.D. studies, the inherently interdisciplinary nature of many of my research projects led me to working with many wonderful people from different research groups. I sincerely thank everybody I have been collaborating with over the course of my studies.

I would like to thank Aya Wallwater and Andrey Bernstein for their contributions to the work on energy harvesting adaptive algorithms. I remain in awe of Aya's and Andrey's mathematical rigor and distinct approaches to problem solving.

I would like to thank John Sarik for his contributions to a wide range of our joint projects – motion energy measurements, examination of student learning, and at least a half-dozen smaller undertakings. I sincerely believe working with John has been the best an interdisciplinary collaboration can possibly be. With John as a "partner in crime", I ventured much further in my research than I would have dared to go on my own.

Special thanks go out to Robert Margolies, with whom we collaborated on a number of EnHANTs project initiatives. Rob's technical expertise, extraordinarily work ethics, and superior organizational skills always kept me reaching for ever-higher standards in our joint work.

Thanks to the entire EnHANTs team! I feel privileged to have had a chance to contribute to such an exciting project. In particular, I would like to thank professors Peter Kinget and John Kymissis for their continuing support. Additionally, it has been a real pleasure to work closely with Baradwaj Vigraham, Jianxun Zhu, Marcin Szczodrak, Tarun Sharma, Zainab Noorbhaiwala, Dongzhen Piao, Luis Pena, Matthew Cowan, Albert Maldonado, Gerald Stanje, Dillen Roggensinger, Haodan Huang, Abraham Skolnik, Enlin Xu, Mina Cong, Sonal Shetkar, Kanghwan Kim, Chang Sun, Michael Zapas, and many other students that dedicated their time to the EnHANTs project. Additionally, I sincerely thank the students who trusted me to be their research project supervisor. I tried my hardest to be the very best supervisor I could be.

I owe a debt of gratitude to my M.Sc. advisor, Dr. Peter Mason, and to my Telcordia Technologies mentor, Dr. Scott Alexander, for giving me the confidence to set out on the adventure of a Columbia University Ph.D. program.

Last but not least, I would like to thank my parents. Their determination and courage continue to inspire me in my most ambitious undertakings.

Support: This thesis was supported in part by a Columbia University Presidential Fellowship, Canadian Graduate Scholar (CGS) NSERC Fellowship, 2012 Google Anita Borg USA Fellowship, NSF grants CNS-0916263, CCF-0964497, and CNS-1054856, and Vodafone Americas Foundation Wireless Innovation Project. To my husband Valentine. Val, you are the source of my strength and the love of my life.

Chapter 1

Introduction

Recent advances in ultra-low-power wireless communications and in energy harvesting will soon enable networks of self-powered wireless devices. Such networks will serve as building blocks for different Internet of Things (IoT) applications. For example, they can be used to enable new tracking and monitoring applications for supply chain management (continuous monitoring of objects, keeping track of object configurations), wearable computing, and smart buildings. With energy harvesting (also known as energy scavenging), nodes can derive energy from environmental sources, such as light and motion [68]. However, in many commonplace environments energy availability is low. For example, the amount of light energy available indoors is a thousand times less than the amount of energy available outdoors [77]. Wireless communication technologies have only recently reached a point at which it is possible to network nodes powered by such low-intensity sources. Correspondingly, numerous challenges need to be addressed before the ubiquitous object networking vision is fully realized. This thesis focuses on the following questions.

- What are the properties of environmental energy sources for ultra-low-power energy harvesting nodes? Powering wireless nodes with ambient energy sources only recently became possible. Correspondingly, there is lack of data and analysis regarding energy availability, variations, and influencing factors.
- How shall nodes adapts their communications and networking to environmental energy conditions? A network of nodes powered by environmental energy sources (such



Figure 1: A network of wireless ultra-low-power energy harvesting networked nodes.

as a network shown schematically in Fig. 1) needs to adapt to energy conditions on all layers of the protocol stack. Furthermore, as environmental energy availability is low, such adaptations require specifically designed low-complexity communication, networking, and resource allocation algorithms.

• How to design ultra-low-power energy harvesting networked nodes? – While there is keen interest in creating small ultra-low-power devices that can be attached to different IoT objects, existing nodes are heavy, bulky, and use orders-of-magnitude more energy than the energy available in commonplace environments. Overcoming these limitations requires cross-layer interactions between energy harvesting, communications, circuit design, and networking.

To outline the contributions of this thesis, in this chapter we first provide the background related to energy sources and hardware for energy harvesting nodes (Section 1.1), introduce the *design space* for energy harvesting adaptive algorithms (Section 1.2), and provide a high-level overview of the ultra-low-power nodes we designed and prototyped – the Energy Harvesting Active Networked Tags (Section 1.3). We overview the contributions of this thesis in Section 1.4.

Section 1.1 is partially based on the material we presented in [23,32]. Section 1.2 is based on the material we presented in [30,31]. Section 1.3 is based on the material we presented in [23,25,32].

1.1 Energy Harvesting and Storage for Ultra-low-power Energy Harvesting Nodes

In traditional (non-energy-harvesting) nodes, all energy a node can use throughout its lifetime is pre-stored in a non-rechargeable battery. Energy harvesting nodes, on the other hand, *obtain their energy from the environment*. In this section, we briefly describe energy harvesting and energy storage technologies for ultra-low-power energy harvesting nodes.

Light energy harvesting – Light is one of the most abundant energy sources, with typical *ir-radiance* (total energy projected and available for collection) ranging from 100 μ W/cm² in indoor environments to 100 mW/cm² in direct sunlight (note the significant difference) [68,77].

Various solar cell technologies are available for light energy harvesting. The energy conversion efficiency of a solar cell is defined as the percentage of the available energy that is harvested by the solar cell. Conventional single crystal and polycrystalline solar cells, such as those commonly used in calculators, have conversion efficiencies of around 10%–20% in direct sunlight. However, their efficiency rapidly declines with a reduction in energy availability (i.e., they are less efficient with dimmer sources). Additionally, conventional solar cells are rigid (inflexible), which makes it difficult to attach them to non-rigid items (e.g., clothing, paperback books). Numerous lightweight mechanically flexible solar cells are becoming commercially available, and new technologies are being actively developed [69]. Efficiencies of these solar cells are typically lower than the efficiencies of conventional single crystal and polycrystalline solar cells. For solar cells based on organic semiconductors [71], for example, the efficiency is typically 1%–1.5%.

Motion energy harvesting – Another potential source of energy is *kinetic (motion) energy*. In industrial scenarios, certain machine motions may correspond to substantial energy availability (e.g., rotating turbines). In commonplace environments, *human activities* such as walking can generate substantial power [86], and harvesting even a small fraction of that power is useful for energy harvesting nodes.

There are several possible ways of harvesting motion energy. For example, in *piezoelectric* harvesters, energy is generated from straining a material. An example of piezoelectric energy harvesting is energy harvesting through footfall, where a harvesting device is placed in a shoe and energy is

generated and captured with each step [51]. In *inertial* harvesters, power is drawn from the relative motion between an oscillating proof mass and the frame from which the mass is suspended [60]. Inertial harvesters can be manufactured in a form factor suitable for the IoT applications. Such harvesters are currently under active development (while large devices, such as motion-based mobile phone chargers [96] and "shake flashlights" have long existed, small form-factor harvesters are not yet commercially available). An inertial harvester suitable for a small IoT device (e.g., under 5 cm x 5 cm, and weighting less than 2 grams) can generate 100–200 μ W from human walking [39,103]. **Other energy sources for energy harvesting** – Energy harvesting nodes may potentially obtain energy from other sources, such as radiowaves and temperature gradients [58,68]. However, these sources are currently of a limited use for powering nodes in commonplace environments. The power available from radiowaves, for example, is *100 times less than the power available from indoor light* [104].

Energy storage – Without the ability to store energy, an energy harvesting node would be able to operate only when directly powered by the environmental energy. For an ultra-low-power node, energy storage components need to be compact and efficient, and need to have very low self-discharge rates.

Rechargeable batteries are an excellent option for energy storage, and numerous battery options are available. Thin film batteries are particularly attractive for ultra-low-power IoT nodes since they can be made flexible. Use of capacitors for storing harvested energy recently started gaining attention [35, 43, 120]. Capacitors can be cycled many more times than batteries. The disadvantage of using capacitors, however, is that large electrolytic capacitors self-discharge over hours or days. The energy density (how much energy can be stored per unit of volume) of capacitors is also much lower. A typical battery can store about 1000 J/cm³, whereas high performance ceramic capacitors can store 1-10 J/cm³.

1.2 Design Space for Energy Harvesting Adaptive Algorithms

Within the overall energy harvesting adaptive networking space, there is a wide variety of different scenarios, which call for different algorithmic approaches. In this section, we briefly introduce several



Figure 2: Examples of different light energy sources: (a) *deterministic profile* (Las Vegas, NV [95], outdoors), (b) *partially predictable profile* (New York City, NY, a stationary indoor node), and (c) *stochastic process* (New York City, NY, a mobile node in Times Square at nighttime).

dimensions of the *algorithm design space* for energy harvesting nodes. The combinations of values along these dimensions induce several "working points", some of which we study in this thesis.

Environmental energy model – The model representing harvested energy depends on various parameters such as the energy source (e.g., solar or kinetic), properties of the environment, and node behavior (e.g., stationary, semi-stationary, or mobile).

Fig. 2 provides examples of light energy sources in different settings. Fig. 2(a) shows the light energy recorded over 4 days by National Renewable Energy Laboratory [95] in Las Vegas, NV. Fig. 2(b) shows the light energy recorded, over the same 4 days, in an *indoor* environment. This data was obtained as part of our light energy measurements campaign (see Chapter 3). In Fig. 2(a), the energy availability is time-dependent and predictable. On the other hand, in Fig. 2(b) the energy is time-dependent and periodic, but harder to predict. Time-dependent and somewhat periodic energy source behaviors (along with other inputs such as weather forecasts) allow to develop an *energy profile* [19, 46]. We will refer to energy profiles that accurately represent the future as *deterministic profiles*, and to those that are inaccurate as *partially predictable profiles*.



Figure 3: An example of harvested power versus energy storage "curves" for a light energy harvesting node that uses a capacitor for energy storage.

Energy behavior that does not warrant a time-dependent profile appears in Fig. 2(c). Fig. 2(c) shows the irradiance recorded by a *mobile device* carried around Times Square in New York City at *nighttime*, obtained as part of our light energy measurements campaign. 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 it is 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.

Energy storage type – As indicated in the previous section, in order to operate when not directly powered by the environmental energy, energy harvesting devices require an *energy storage* component – 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 [78]. For *capacitors*, we consider the *nonlinearity* of capacitor-based energy harvesting 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* [35, 56].¹ The nonlinear relations are demonstrated in Fig. 3.

Energy storage capacity versus amount of energy harvested – Energy storage capacity can vary from 0.16 J for an EnerChips solid state energy storage device [90] to 4700 J for an AA battery. The environmental energy availability also varies widely, from thousands of $J/cm^2/day$ in sunny outdoor conditions to under 2 $J/cm^2/day$ in indoor environments (see Chapter 3).

Time granularity – Nodes can characterize the harvested energy and make decisions on timescales

¹Solar cells have highly nonlinear output-versus-voltage characteristics. In simple energy harvesting 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.



Figure 4: The envisioned EnHANTs: (a) the intended EnHANT form factor, (b) EnHANTs in comparison to sensor networks and RFIDs, and (c) an organic semiconductor-based *mechanically flexible* solar cell developed in the Columbia Laboratory for Unconventional Electronics.

from seconds to days. This timescale is related to the storage-to-energy-harvesting ratio and to the environmental energy model.

Problem/network size – Energy harvesting affects nodes' individual decisions, pairwise (link) decisions, and behavior of networked nodes (e.g., flow control, topology determination, routing).

1.3 Energy Harvesting Active Networked Tags (EnHANTs)

One of the challenges considered in this thesis is the design of ultra-low-power energy harvesting nodes. Specifically, we focus on the design of *Energy Harvesting Active Networked Tags (EnHANTs)*, which will be a new class of ultra-low-power wireless nodes. The design of the EnHANTs is influenced by the properties of environmental energy sources we study in this thesis. Moreover, the EnHANTs serve as a platform for experimenting with energy-adaptive algorithms we develop in this thesis for energy harvesting nodes.

The envisioned EnHANTs will be devices that:

- **Network** Actively communicate with one another and with EnHANT-friendly devices in order to forward information over a multihop network.
- Operate at ultra-low-power Spend a few nJ or less on every communicated bit.
- Harvest environmental energy Collect and store energy from sources such as light and motion.
- Are energy adaptive Adapt communications and networking to satisfy energy harvesting constraints.

- Exchange small messages Exchange limited information (i.e., mostly IDs) using low data rates.
- **Transmit to short ranges** Communicate only when in close proximity (1 to 10 meters) to one another.
- Are thin, *mechanically flexible*, and small A few square centimeters at most.

The envisioned EnHANT form factor is shown in Fig. 4(a). In terms of complexity, throughput, size, and energy requirements, EnHANTs fit between *RFIDs* and *sensor networks*, as shown schematically in Fig. 4(b). Similarly to RFIDs, EnHANTs can be attached to commonplace objects. Presence of power sources (via energy harvesting) and distributed multihop operation shift EnHANTs closer to sensor networks than to RFIDs that are traditionally passive and not networked. Compared to sensor nodes, however, EnHANTs operate at significantly lower data rates, consume less energy, and transmit mostly ID information.

EnHANTs will support a variety of tracking and monitoring applications beyond what RFIDs permit. While RFIDs make it possible to *identify* an object in proximity to a reader, EnHANTs will make it possible to *search for an object in a network of devices*, and to continuously track objects' whereabouts and their proximity to each other. One application that demonstrates the building blocks of an EnHANTs-based application is *misplaced library book locator*. In this application, library books will be able to identify those among themselves that are significantly misplaced (e.g., in an incorrect section), and report the misplacement. To accomplish this task, each book is assigned a unique ID using an assignment scheme closely related to the Dewey Decimal Classification. Each book has a light-powered tag (EnHANT) that can transmit and receive information within a radius of one meter or less and can perform some basic processing. Nearby books wirelessly exchange IDs. The IDs of books that appear out of place are further forwarded through the network of books, eventually propagating to sink nodes.

The same building blocks used in the library applications can enable several other applications. A variety of items can be tracked, and a range of possible desirable or undesirable object configurations can be queried for, and can trigger reports. Examples include finding items with particular characteristics in a store, continuous peer monitoring of merchandize in transit, locating misplaced items (e.g., keys or eyeglasses), and locating survivors in disasters such as structural collapse.

EnHANTs can potentially be implemented by combining organic solar cells with communications chips supporting Ultra-Wideband Impulse-Radio (UWB-IR):

Energy harvesting with organic solar cells – Advances in the area of organic semiconductors for energy harvesting allow the fabrication of organic solar cells on *flexible substrates* [71], thereby enabling pervasive use of future mechanically flexible EnHANTs. An array of flexible solar cells recently designed in the Columbia University Laboratory for Unconventional Electronics is shown in Fig. 4(c).

Ultra-low-power ultra-wideband communications – UWB-IR is a compelling technology for short range ultra-low-power wireless communications [15, 16, 108]. It uses very short pulses (on the order of nano-seconds) that are transmitted at regular time intervals, with the data encoded in the pulse amplitude, phase, frequency, or position. At low data rates, the short duration of the pulses allows most circuitry in the transmitter or receiver to shut down between the pulses, resulting in significant power savings compared to traditional narrow-band communication systems.

Practical CMOS IR circuits with energy consumption on the order of a nano Joule per bit have recently been demonstrated. For example, [14] demonstrated a UWB-IR receiver and transmitter that require 1.65 nJ/bit and 280 pJ/bit, respectively, at a data rate of 1 Mbit/s. Research indicates that UWB-IR transceivers in the 3–5 GHz band with data rates of 0.1–1 Mbit/s and receiver and transmitter consumption of less than 500 pJ/bit and 50 pJ/bit, respectively, are within reach.

To put these numbers in perspective, consider an EnHANT with a 10 cm² organic semiconductor solar cell. Outdoors, the system would harvest $10 \text{cm}^2 \cdot 100 \text{mW/cm}^2 \cdot 0.01 = 10 \text{ mW}$. Under the assumption that receiving a bit requires 1 nJ, the achievable data rate would be $(10 \cdot 10^{-3})/(1 \cdot 10^{-9}) = 10 \text{ Mbit/s}$. The achievable data rate in indoor environments would be $(10^{-5})/(1 \cdot 10^{-9}) = 10 \text{ Kb/s}$. Such data rates should be sufficient for EnHANTs to communicate and network.

1.4 Summary of Contributions

This thesis focuses on *ultra-low-power energy harvesting networked nodes*. In this general area, we made contributions to energy source characterization, algorithm design, and node design and

prototyping. We also made contributions to project-based engineering education. In this section, we briefly summarize the contributions of this thesis.

1.4.1 Characterizing Environmental Energy for Energy Harvesting Nodes

In Chapters 3 and 4, we present characterizations of environmental energy availability for energy harvesting IoT nodes. Until recently, harvesting low levels of ambient energy was impractical. Correspondingly, few efforts to characterize ambient energy sources were undertaken. In this work we determined energy availability and properties for *indoor light* and *human and object motion*. Our insights are important for system design (e.g., determining harvester size, battery size) and algorithm design (e.g., defining available data rates, expected energy variations, predictability). The specific contributions are summarized below.

Light energy harvesting (Chapter 3): To characterize indoor light energy, we conducted a *first-ofits-kind long-term indoor irradiance (light energy) measurement study.* We designed and developed a system for long-term irradiance data collection, and deployed the developed systems in a set of locations in Columbia University for over 1.5 years. Using the collected *light energy traces*, we obtained insights into indoor light energy availability and characteristics. The characterizations demonstrated the feasibility of powering IoT nodes using indoor light energy harvesting. They also demonstrated that light energy availability levels are substantially different even within seemingly uniform environments, and that simple parameters can substantially improve light energy predictions. The collected traces can be used as inputs to energy harvesting system simulators and emulators.

We summarized the findings of the study in [30, 31], and shared the indoor light energy traces with the community on enhants.ee.columbia.edu and via CRAWDAD at [33].

Kinetic (motion) energy harvesting (Chapter 4): Characterizing motion energy is more complex than characterizing light energy because it requires in-depth characterizations of motion frequencies and amplitudes. We examined energy corresponding to moving objects, specific human motions (walking, running, cycling), and daily human routines. For specific human motions, we used a dataset with over 40 participants [111], obtaining extensive and general human motion kinetic energy characterizations. To characterize motion energy associated with daily human routines, we conducted an acceleration measurements campaign with 5 participants over a total of 25 days, collecting traces with more than 200 hours of acceleration information. The results indicated that energy availability associated with object motion is low, and demonstrated the range of motion frequencies and harvested powers for different participants and activities. The results also highlighted the importance of human physical parameters for energy harvesting, and demonstrated that the power generation process associated with human motion is highly variable, with only brief intervals of high power levels.

1.4.2 Resource Allocation Algorithms for Energy Harvesting Nodes

In Chapter 5, we summarize the modeling frameworks and the design, development, and performance evaluations of resource allocation algorithms for ultra-low-power energy harvesting nodes. The energy available to energy harvesting nodes varies in time. Motivated by the needs of tracking and monitoring IoT applications to communicate consistently, we aimed to allocate energy harvesting nodes' resources *in a uniform way with respect to time*. To formulate the energy allocation problems, we used the *utility maximization* and the *lexicographic maximization* frameworks. These frameworks are typically applied to achieve *fairness among nodes* [7, 19, 48, 55, 67]. We applied them to achieve *time-fair* resource allocations.

We mainly considered *deterministic energy profile* and *stochastic* environmental energy models, for battery-based systems and for capacitor-based systems, and focused on the cases of a *single node* and a *node pair (link)*. For the *deterministic profile energy model*, we obtained data rate and energy spending rate allocations (once the energy spending rates are determined, they can be converted to duty cycles, sensing rates, or communication rates). We examined optimal policies, as well as a number of simple policies, for which we obtained performance guarantees. For the *stochastic energy model*, we considered the case in which the energy inputs are i.i.d. random variables, and formulated allocation problems as *average-cost Markov Decision Processes* (MDPs). We examined an approximation to the optimal policy that uses the *uniform discretization of the problem*, and obtained a bound on the performance degradation due to discretization. We additionally examined several simple policies, for some of which we provided performance guarantees.

We summarized some of the contributions in [21, 30, 31].

1.4.3 EnHANTs Prototypes Testbed for Energy Harvesting Adaptive Policy Evaluations

In Chapter 6, we focus on the design and the prototyping effort for Energy Harvesting Active Networked Tags (EnHANTs) and the EnHANTs prototype testbed. Features of the designed EnHANTs prototypes include energy harvesting using organic solar cells, multihop data forwarding over UWB-IR physical layer, real-time energy storage state tracking, and adaptations of communications and networking to the environmental energy. Specifically, the prototypes implement energy harvesting adaptive energy spending rate control, flow control, and topology selection functionalities. The designed EnHANTs testbed uniquely incorporates a software-based light energy control setup, which allows evaluating algorithms in controllable environments with real energy harvesting hardware.

We presented the high-level EnHANTs design in [23,32]. The EnHANTs prototype and testbed design and development were conducted over 4 years in 6 iterative phases, each of which was presented in a conference demonstration session [26, 28, 57, 80, 85, 119]. We summarized the EnHANTs prototype and testbed design in [25].

We used the testbed to evaluate energy harvesting adaptive algorithms (i.e., the contributions we describe in Chapter 5) with the light energy traces (i.e., the contributions we describe in Chapter 3). To the best of our knowledge, our work is the first attempt to evaluate energy harvesting adaptive policies in a controllable experimental environment.

1.4.4 Project-based Learning

Appendix A summarizes our contributions to engineering education. Modern engineering landscape will increasingly require system engineering skills that are not typically acquired in traditional engineering and computer science programs. Thus, over 11 semesters, we engaged more than 50 high school, undergraduate, and Masters students in more than 100 interdisciplinary research projects related to the energy source characterizations and to the EnHANTs prototype and testbed design and development. To the best of our knowledge, our experience with organizing multiple student projects to contribute to a large-scale effort is unique. We describe the challenges and the solutions associated with our approaches to project-based learning, and present the results of a survey-based

assessment that demonstrate the effectiveness of our approaches.

The contributions presented in Appendix A were summarized in [27].

Chapter 2

Related Work

In this chapter, we provide a brief overview of the related work. Energy efficiency in wireless networks has long been a subject of research (see reviews [3, 45, 59]). In comparison, networking *energy* harvesting nodes has only recently started gaining attention. We first overview the work related to energy source characterization for energy harvesting nodes (Section 2.1). Then, we provide an overview of the work related to energy harvesting adaptive algorithms (Section 2.2). We then describe the work related to ultra-low-power energy harvesting node design and prototyping (Section 2.3). Finally, we overview the work related to student project organization (Section 2.4).

2.1 Energy Source Characterization for Ultra-low-power Energy Harvesting Nodes

In this work, we characterize *indoor light energy* and *kinetic (motion) energy* for ultra-low-power energy harvesting nodes.

Light energy – Since large-scale *outdoor* solar panels have been in used for decades, properties of the Sun's energy were examined in depth [50,77,95]. Practical outdoor solar energy considerations for energy harvesting in sensor networks (e.g., light obstructions, scattering) were discussed in [88]. 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 [36,79]. 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 [35,94]. Photometric measurements, however, do not provide energetic characterization, and there is lack of data (e.g., traces) and analysis (e.g., variability, predictability, and correlations) regarding energy availability [77].

Kinetic energy – In this work, we characterize energy of object and human motion, short-term (i.e., per-activity) and longer-term (i.e., on a scale of days). To the best of our knowledge, the energy availability from *object motion* has not been characterized before. Previous *human motion* energy harvesting studies had a limited number of participants (10 in [39] and 8 in [103,116]) and focused on walking and running on a treadmill at a constant pace. Due to the small sample sizes and number of activities considered, these studies do not lead to general conclusions about energy availability and characteristics for wireless and mobile nodes. To the best of our knowledge, the 40-participant dataset [111] that we analyze is the first publicly available acceleration dataset collected from a large number of participants. This dataset *was not previously used for an energy harvesting study*. A relatively small set of *day-scale human motion* acceleration traces was studied in [116], which focused on determining node energy budgets under assumptions corresponding to relatively large electronic devices. In this work, we collect day-scale data that in some cases has more information per participant, examine the traces under assumptions corresponding to ultra-low-power energy harvesting nodes, and characterize various properties that have not been considered before.

2.2 Energy Harvesting Adaptive Algorithms

This thesis deals with *resource allocation in energy harvesting devices*. The design of energy harvesting-adaptive communication, networking, and resource allocation algorithms has recently been gaining attention. Related work in this area can be classified according to the *environmental energy model* we introduced in Section 1.2:

• Deterministic profile – In [38, 46], duty cycle adaptations (mostly for *single nodes*) are considered. Transmission power adaptation and transmission scheduling for small scenarios (nodes, links) are examined in [89, 112, 113] and [2], respectively. For a *network*, various metrics are considered including data collection rates [19], data retrieval rates [114], throughput maximization [10], and routing efficiency [54].

- **Partially predictable profile** While considering energy predictable, [10, 46, 55, 65] have provisions for adjustments in cases in which the predictions are inaccurate.
- Stochastic process Dynamic activation of energy harvesting sensors is described in [42] for a *single node*, and for a *cluster* in [47]. Admission and power allocation control policies are developed in [20]. Routing and scheduling policies are developed in [49]. Maximizing the utility of the average data rates via joint power allocation and energy management is examined in [40]. Energy allocation policies for source-channel coding are developed in [9].
- Model-free approach Duty cycle adjustments for a single node are examined in [102]. A capacitor-based system is presented and the capacitor leakage is studied in [120].

In this thesis, we examine policies that are aimed at allocating the nodes' dynamic and timedependent resources in a *uniform way with respect to time*. This goal is motivated by the needs of tracking and monitoring systems to communicate *consistently*. The need for policies that enable such behavior in energy harvesting devices has been noted by many researchers [19,46,64,102]. "Smoothing" node duty cycles using a control theory approach is examined in [102]. Energy allocation vectors with minimal variance are sought in [64]. We note that the approach introduced in [117] for throughput optimization in quality-of-service-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 communications buffer constraint [112]). A throughput optimization framework for energy-harvesting nodes [10], 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. *To the best of our knowledge, our approach to resource allocation problem formulations is the first that allows obtaining resource allocations for systems with non-linear energy storage* (i.e., capacitor-based systems).

A number of works on energy harvesting adaptive algorithms pursue different objectives and *lead* to different resource allocations. For example, [89] aims to maximize the short-term throughput, [113] minimizes the packet transmission time, and [11,20,40] maximize the sum of utilities of nodes' timeaveraged communication rates. Achieving these objectives does not result in resource allocations that are uniform with respect to time.

In this thesis we focus on optimal and heuristic policies for the deterministic profile and stochastic environmental energy models, and evaluate the policies both analytically and experimentally. Some of the policies we examine for the *deterministic profile* energy model have been used as heuristics before (for example, in [19,55]). We demonstrate the bounds on the performance of these policies. For the *stationary stochastic* energy model, in this work we use a Markov Decision Process-based (MDP-based) approach to node resource allocation. While MDP-based approaches have been proposed before [47,105], such approaches typically assume that the node's harvested energy and energy storage are discretized.¹ To the best of our knowledge, our work is *the first to examine the performance degradation associated with the energy storage state discretization*. This thesis also uniquely considers *analytically* a set of simple policies for the stochastic energy model that have been used as heuristics before (for example, in [47,63]).

We note that resource allocation in energy harvesting nodes has some similarities with power consumption scheduling in power networks (e.g., [53] 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.

2.3 Designing and Prototyping Energy Harvesting Active Networked Tags

While the idea of pervasive *networks of objects* has been proposed before (e.g., in the Smart Dust project [106]), the harvesting and communications technologies have only recently reached a point where networked energetically self-reliant tags are becoming practical [76,92]. Correspondingly, combining the advances in energy harvesting and ultra-low-power communications has recently attracted increasing attention from industry and academia [91, 100].

The Energy Harvesting Active Network Tag prototypes we have designed and developed offer ¹The energy storage state discretization is assumed in many non-MDP-based approaches as well, e.g., [20, 40, 83]. several advantages over the state of the art. The prototypes include an energy harvesting module which provides real time energy awareness, and organic photovoltaics, which are specifically designed for indoor light energy harvesting. Other existing energy harvesting modules [74, 93, 97] offer only limited energy awareness and do not provide real time harvesting rate information. Existing sensor network nodes that harvest energy from sunlight [46, 88, 114] and indoor light [115] typically use monocrystalline or amorphous silicon solar cells, rather than organic photovoltaics. The EnHANTs prototypes include an ultra-wideband impulse radio (UWB-IR) communication module and support medium access control and networking functionalities over the UWB-IR physical layer. While there have been other UWB-IR implementations [16, 84, 108], to the best of our knowledge, none have implemented and tested functionality above physical layer over the UWB-IR transceivers.

To the best of our knowledge, the EnHANTs testbed we have developed is the first to allow experiments with *real energy harvesting hardware under repeatable controllable energy inputs*. Several energy harvesting testbeds exist (e.g., [18,88,114]) and a few are under development (e.g., [17,44]). These testbeds do not control the amount of energy nodes can harvest. Solar simulators (used for testing solar cells) can provide precisely controlled illumination (approximating sunlight), but cannot create trace-based dynamic light environments for the nodes.

Using the EnHANTs prototypes and the EnHANTs testbed, we evaluated several energy harvesting adaptive policies. While many policies have been proposed (see Section 2.2), the majority of the policies were only evaluated via simulations.

2.4 Project-based Learning

In this thesis, we describe a method for engaging students in project-based learning within a largescale multidisciplinary research effort. Previous research has described structuring student research experiences as a course [34,72], and examined a research-based program aimed at undergraduates [70] and a framework for accommodating undergraduate students in a research group [6,107]. Researchers have also examined methods for providing students with interdisciplinary research opportunities [75] and for increasing student communication and collaborative skills [13]. The necessity of engaging students in large-scale system development projects has been recognized as an important educational objective, and some tools for emulating the development scale have been proposed [81]. However, to the best of our knowledge, our experience with providing many students interdisciplinary projectbased research opportunities as part of a single large-scale ongoing research effort is unique.

Chapter 3

Characterizing Light Energy for Energy Harvesting Nodes

In this chapter, we focus on *light energy* availability and properties for ultra-low-power energy harvesting nodes. Specifically, we consider light energy available *indoors* and light energy available to *mobile nodes*.

To characterize this energy, we conducted a first-of-its kind 16 month-long indoor radiant energy measurements campaign and a mobile outdoor light energy study. Based on our measurements, we obtained insights into energy availability in different environments, and made observations regarding energy variability, predictability, and influencing factors. Our long-term energy measurements campaign resulted in unique *irradiance traces* that can be used as energy inputs to energy harvesting adaptive system and algorithm simulators and emulators. We made the traces publicly available at enhants.ee.columbia.edu, and also published them via the CRAWDAD repository [33]. The insights and observations we obtain are important for designing energy harvesting systems and energy harvesting adaptive algorithms.

Some elements of the presented study were completed as part of undergraduate and M.S. projects (see Appendix A). Specifically, projects of M. Bahlke, M. Zapas, and E. Xu contributed to measurement setup design and development; some of the measurements were carried out by M. Zapas, S. Shetkar, C. Sun, and K. Kim.


Figure 5: A schematic diagram of the relationships between energy parameters: irradiance, I, irradiation, H, energy available to a node, D, and energy harvested by the node, Q.

In this chapter, we first describe our light energy study methodology (Section 3.1). We then comment on *spatial variability* of indoor light energy levels (Section 3.2), *energy budgets* for indoor light energy harvesting nodes (Section 3.3), and on properties of corresponding *environmental energy profiles* (Section 3.4). Additionally, we also briefly discuss our *mobile measurements* and light energy properties for mobile energy harvesting nodes (Section 3.5).

This chapter's contributions were previously presented in [30, 31].

3.1 Methodology

The relationships between variables characterizing light energy availability for energy harvesting systems are illustrated in Fig. 5. Our measurements capture *irradiance*, radiant energy incident onto surface (in W/cm²), denoted by *I*. *Irradiation* H_T (in J/cm²) 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 the given physical properties* (size, efficiency) can harvest in a time slot *i* is denoted by *D*. For a solar cell with area *A* and energy conversion efficiency η , $D = A \cdot \eta \cdot H$. For the energy availability calculations presented in this chapter, we use $A = 10 \text{ cm}^2$ (i.e., a small node with a form factor similar to the future EnHANTs demonstrated in Fig. 4(a)) and $\eta = 1\%$ (i.e., the efficiency of an organic solar cell).

To characterize indoor energy availability, over the period of June 2009 – September 2010 we conducted a light measurement study in office buildings in New York City. In this study we examined a set of short-term indoor and outdoor irradiance measurements, including measurements with *mobile devices*. We also collected a set of *long-term measurements* in several indoor locations. Table 1



Figure 6: A schematic diagram of the indoor irradiance measurements locations L-1–L-6.



Figure 7: Sample irradiance measurements in indoor locations L-2, L-3 and in outdoor location O-1 (Mar. 2, 2010 – Mar. 12, 2010).

provides a summary of the indoor measurement locations. The locations are shown schematically in Fig. 6. For comparison, in addition to our own indoor measurements, we also analyzed a set of outdoor irradiance traces provided by the US Department of Energy National Renewable Energy Laboratory (NREL) [95]. Sample irradiance measurements (for two our setups and one setup from [95] over the same 10 days) are provided in Fig. 7.

For the irradiance measurements we used TAOS TSL230rd photometric sensors [99] 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.

3.2 Spatial Variability

In this section, we demonstrate spatial variability of light energy available for indoor energy harvesting devices. We demonstrate the spatial variability in two sets of environments: on a bookshelf and on a human body.



Figure 8: Irradiance measurements in different parts of a bookshelf: (a) the floor plan of the office where the measurements were taken, (b) the leftmost bookshelf stall (stall A) with measurement locations identified, and (c) the irradiance values in different locations on the bookshelf (measurements taken when outdoor light did not affect the values).



Figure 9: Irradiance measurements I (in $\mu W/cm^2$) in different locations on a person.

- **Bookshelf** We measured the irradiance levels at different locations on a bookshelf which consists of 4 stalls with 7 shelves in each, and is illuminated by overhead fluorescent lights as shown in Fig. 8(a). We took measurements at the intersection points of shelves and bookshelf stalls, as depicted in Fig. 8(b). Fig. 8(c) shows the measured irradiance levels for each of the shelves. It can be seen that the irradiance levels vary as a function of the distance (and orientation) from the light sources.
- Human body We additionally measured the irradiance levels at different locations on a person, for a person standing directly under an overhead lamp. The corresponding "map" of the irradiance values (in μW/cm²) recorded in different location on a human body is shown in Fig. 9. We note that there is a 9.8 times difference in the irradiance levels between the different parts of the human body.

Location	Location description	Experiment	H_d	$\sigma(H_d)$	r
index		timeline	$(J/cm^2/day)$		(Kb/s, cont.)
L-1	Students' office. South-facing, 6th floor	Aug. 15, 2009 –	1.3	0.72	1.5
	above ground, windowsill-located setup.	Sept. 13, 2010			
L-2	Students' office (same office as setup	Nov. 13, 2009 –	1.28	0.76	1.5
	L-1). Setup on a bookshelf far from win-	Sept. 9, 2010			
	dows, receiving direct sunlight for a short				
	portion of a day.				
L-3	Departmental conference room. North-	Nov. 7, 2009 -	63.0	48.0	72.0
	facing, 13th floor above ground,	Sept. 13, 2010			
	with large unobstructed windows.				
	Windowsill-located setup.				
L-4	Students' office. South-West-facing,	Nov. 5, 2009 $-$	9.2	6.9	7.9
	windowsill-located setup.	Sept. 29, 2010			
L-5	Students' office (directly under the office	June 25, 2009 –	12.3	8.3	13.9
	of setup L-1). Windowsill-located setup.	Oct. 11, 2009			
L-6	Students' office. East-facing, often re-	Feb. 15, 2010 $-$	97.3	64.4	112.3
	ceiving unattenuated reflected outdoor	Sept. 20, 2010			
	light. Windowsill-located setup.				
0-1	Outdoor: ECSU meteostation [95], Eliz-	Jan. 1, 2009 $-$	1517	787	1,750
	abeth City, NC.	Dec. 31, 2009			
O-2	Outdoor: HSU meteostation [95], Ar-	Jan. 1, 2009 $-$	1407	773	1,600
	cata, CA.	Dec. 31, 2009			

Table 1: Light energy measurement setups – average daily irradiation and achievable bit rates.

These results demonstrate that the available energy levels are substantially different even within seemingly uniform environments.

3.3 Node Energy Budgets and Daily Energy Availability

One use of the measurements we have collected is to determine *energy budgets* for indoor light energy harvesting nodes. We calculate the total daily *irradiation* H_d , representing energy incident onto 1 cm² area over the entire course of a day. For example, for setup L-1, the H_d values are shown in Fig. 10. Table 1 presents the average and the standard deviation values, correspondingly $\overline{H_d}$ and $\sigma(H_d)$. 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 use of shading, windows, and indoor lights.

Table 1 also shows the bit rate r a node would be able to maintain throughout a day when exposed to the irradiation H_d . As the energy cost to communicate, we use 1 nJ/bit, which corresponds to



Figure 10: Long-term daily irradiation, H_d , for setup L-1 (Aug. 15, 2009 – Sept. 13, 2010). energy spending of ultra-low-power transceivers envisioned for the EnHANTs, as we described in Section 1.3. Specifically, we calculate the bit rate r as $r = A \cdot \eta \cdot H_d/(3600 \cdot 24)/(10^{-9})$. These bit rates can be seen as the "communication budgets" for light energy harvesting devices (such as EnHANTs) deployed in indoor environments.

To predict the daily energy availability H_d , a node can use a simple exponential smoothing approach (also known as exponential averaging), calculating a predictor for a 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 [50,82]. We examined whether the H_d values in the indoor settings are correlated with the weather data provided in [101]. We determined statistically significant correlations for all setups except L-2.¹ 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 patterns also influence indoor radiant energy in indoor 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 $0.26\overline{H_d}$.

¹The correlation coefficients of the H_d values with the weather data are as follows: L-1: $\rho = 0.35$ (p < 0.001), L-2: no statistically significant (p < 0.05) correlation, L-3: $\rho = 0.137$ (p < 0.05), L-4: $\rho = 0.29$ (p < 0.001), L-5: $\rho = 0.24$ (p < 0.05), L-6: $\rho = 0.71$ (p < 0.001).



Figure 11: 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 $\sigma(H_T)$.

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, $\rho = 0.56$ (p < 0.001), and for setups L-1 and L-5 facing in the same direction, $\rho = 0.71$ (p < 0.001).² 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.

²We also detected the following statistically significant correlations: {L-1, L-3}: $\rho = -0.19 \ (p < 0.05)$, {L-3, L-4}: $\rho = 0.52 \ (p < 0.001)$, {L-3, L-6}: $\rho = 0.25 \ (p < 0.05)$, {L-4, L-6}: $\rho = 0.47 \ (p < 0.001)$.

3.4 Energy Profiles

To characterize energy availability at different times of day, we determine the H_T values for different 0.5 hour intervals T, thus generating *energy profiles* for the different locations. Sample energy profiles are shown in Fig. 11, 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 substantially different. For example, while setup L-3 exhibits daylight-dependent variations in irradiance, for setup L-2 the irradiance is either 0 or 45 μ W/cm² for most of the day (as this setup is illuminated mostly by indoor light). In addition, while for setup L-2 the lights are often turned 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 [1,50]. Specifically, 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, $H_T(j)$. 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, $\sum_{i>j} H_T(i)$.³ We determined that such correlations are present in indoor environments, but they are generally weaker than in outdoor settings. For example, for the outdoor setup O-1 the correlation between the energy received in the 10:30–11:00 AM time slot and in the 16:30–17:00 PM time slot is $\rho = 0.5$ (p < 0.001), while for the indoor setup L-1 it is $\rho = 0.2$ (p < 0.001). For the outdoor setup O-1, the correlation between the amount of energy received before 08:00 AM and the amount of energy received after 08:00 AM is $\rho = 0.77$ (p < 0.001), while for the indoor setup L-3 it is $\rho = 0.31$ (p < 0.001). These results suggest that energy profile prediction techniques developed for outdoor energy harvesting environments (such as [1,50]) may be extended to indoor environments, but their performance indoors is likely to be worse.

³Additional correlation results are available in [29].

-	0 00	0	-			
Meas.	Measurement description	Experiment	Experiment	\overline{I} ,	$\sigma(I)$	r
index		start time	duration	$(\mu W/cm^2)$		(Kb/s)
LM-1	Pedestrian walking around university campus (indoor and outdoor environ-	4/5/2010, 13:06	1:00:09	4,700	9,760	470
LM-2	ments) carrying a measurement setup. Commuting on public transit, measure- ment setup attached to a backpack, mea- surements outdoors, indoors (subway, train, office).	7/13/2010, 15:02	1:40:30	134	448	0.45
LM-3	Car-based roadtrip, measurement setup attached to the dashboard.	7/16/2010, 12:26	2:57:00	11,416	4,370	38.1
LM-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
LM-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
LM-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

Table 2: Mobile light energy measurements – average irradiance and achievable bit rates.



Figure 12: Irradiance measurements recorded by a mobile device: a mix of indoor and outdoor conditions (note the log scale of the y-axis).

3.5 Mobile Measurements

We also conducted shorter-term experiments for *mobile devices*. Table 2 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 LM-6) was shown in Fig. 2(c). Fig. 12 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 LM-1). The measurements demonstrated in Fig. 12 highlight the disparity between the light energy available indoors and outdoors. For example, inside a lab, the irradiance was 70 μ W/cm², while in sunny outdoor conditions it was 32 mW/cm². 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.

3.6 Conclusions

In this chapter, we summarized our characterization of *indoor light energy availability* for ultralow-power energy harvesting nodes. Our characterizations, based on a first *long-term indoor light energy measurements campaign*, demonstrate the feasibility of powering ultra-low-power nodes using indoor light energy harvesting, and provide insights into the design of energy harvesting nodes and algorithms for such nodes. For example, we demonstrate that the light energy availability levels are substantially different even within seemingly uniform environments, and that in indoor environments the energy models are mostly partially predictable. We also show that simple parameters can significantly improve light energy predictions.

In later chapters, we use the indoor light energy traces we described in this chapter to provide energy inputs for energy harvesting adaptive nodes, in order to obtain numerical results for energy harvesting adaptive policies (Chapter 5) and to examine the performance of energy harvesting adaptive nodes in an energy harvesting testbed (Chapter 6).

Chapter 4

Characterizing Kinetic Energy for Energy Harvesting Nodes

In this chapter, we focus on availability and properties of *kinetic (motion)* energy for ultra-lowpower energy harvesting nodes. Characterizing kinetic energy of unrestricted naturally occurring motions is a complex task [60]. In particular, since *human* motion is a combination of low frequency vibrations that vary from activity to activity and from person to person [8,39,103], characterizing kinetic energy that corresponds to human motion requires in-depth study of motion properties (e.g., the frequencies associated with different types of motions) and human mobility patterns.

Our motion energy study is based on object and human motion acceleration traces that we collected and traces collected for over 40 participants in a dataset [111] (the traces in the dataset [111] were collected for an activity recognition study and have not been analyzed from energy harvesting considerations before). Previous human motion energy harvesting studies only obtained "rule of thumb" estimates for a small set of activities based on small numbers of participants [8, 39, 103]. Using the dataset [111], we obtain extensive and general kinetic energy characterization for a set of common human motions (e.g., walking, running, cycling). Additionally, to study the energy generation processes associated with day-scale human routines, we conducted an acceleration measurements campaign with 5 participants, collecting traces with over 200 hours of acceleration information. Our study demonstrates the range of motion frequencies and harvested powers for



Figure 13: (a) A second-order mass-spring system model of a harvester with proof mass m, proof mass displacement limit Z_L , spring constant k, and damping factor b, and (b) the frequency response magnitude for two different harvesters, H1 and H2.

different motions, participants and activities, shows the importance of human physical parameters for energy harvesting, and provides insights that are important for algorithm design.

The methodology of the study is based on joint work with J. Sarik. Some elements of the study were completed as part of undergraduate and M.S. projects (see Appendix A). Specifically, M. Cong partially examined the dataset [111] and collected several measurements. S. Shetkar contributed to the development of measurement setup and study methodology.

In this chapter, we first describe the harvester model and our measurement setup (Section 4.1) and briefly comment on the energy associated with object motion (Section 4.2). We then describe motion energy availability and characteristics for common human motions (Section 4.3) and for day-scale human routines (Section 4.4).

4.1 Model and Measurement Setup

Our motion energy study is based on *acceleration traces*, which are processed to determine the energy generated by an inertial harvester. In this section, we describe the harvester model and the collection of acceleration measurements. We also detail the procedure for obtaining the power generated by the harvester from the recorded traces and the procedures for determining the harvester parameters.

4.1.1 Inertial Harvester Model

An inertial harvester can be modeled as a second-order mass-spring system with a harvester proof mass m, proof mass displacement limit Z_L , spring constant k, and spring damping factor b [61,103, 118]. Fig. 13(a) demonstrates such harvester model.

Two important harvester design parameters are m and Z_L . The harvester output power, P, increases linearly with m [5], and is non-decreasing (but generally non-linear) in Z_L . Yet, m and Z_L are limited by the harvester weight and size considerations, which ultimately depend on the application. We use the following values that are consistent with the mobile systems' restrictions on the size and weight of the overall node, and correspond to one of the harvester configurations examined in [103]:

- Harvester proof mass, $m = 1 \cdot 10^{-3}$ kg (1 gram), and
- Harvester proof mass displacement limit, $Z_L = 10$ mm.

The other two model parameters, k and b, are tuned to optimize the energy harvested for given motion properties. The parameter k determines the harvester resonant frequency,

$$f_r = 2\pi \sqrt{\frac{k}{m}}.$$

To maximize power output, the resonant frequency, f_r , should match, reasonably closely, the dominant frequency of motion, f_m .

Jointly, k and b determine the harvester quality factor,

$$Q = \frac{\sqrt{k \cdot m}}{b},$$

which determines the spectral width of the harvester. A harvester with a small Q harvests a wide range of frequencies with a low peak value, while a harvester with a large Q is finely tuned to its resonant frequency f_r . The role of f_r and Q can be observed in Fig. 13(b), which shows the magnitude of the frequency response of two different harvesters, denoted by H1 and H2. For H1, $f_r = 2.06$ Hz (which corresponds to a typical frequency of human walking) and Q = 2.35 (k = 0.17, b = 0.0055). For H2, $f_r = 2.77$ Hz (which corresponds to a typical frequency of human running) and Q = 3.87 (k = 0.30, b = 0.0045).



Figure 14: Acceleration measurement unit and placements: (a) our sensing unit which is based on a SparkFun ADXL345 evaluation board, and (b) the sensing unit placements in a multi-participant human motion characterization study [111].

4.1.2 Collecting Motion Information

In this thesis, we examine measurements that we collected and measurements provided in an acceleration dataset of common human motions [111]. Our measurements were obtained using sensing units based on SparkFun ADXL345 evaluation boards (see Fig. 14(a)). Each unit includes an ADXL345 tri-axis accelerometer, an Atmega328P microcontroller, and a microSD card for data logging.¹

The sensing units record acceleration along the x, y, and z axes, $a_x(t)$, $a_y(t)$, $a_z(t)$, with a 100 Hz sampling frequency. We conducted multiple experiments with different sensing unit placements, as we will describe in Sections 4.2 and 4.4.

The dataset of [111] was obtained using an ADXL330 tri-axis accelerometers with a 100 Hz sampling frequency. The measurements of [111] were conducted with sensing unit placements corresponding to a shirt pocket, waist belt, and trouser pocket, as shown in Fig. 14(b).

In all the measurements, the orientation of the sensing unit was not controlled. We examine the overall magnitude of acceleration, $a(t) = \sqrt{a_x(t)^2 + a_y(t)^2 + a_z(t)^2}$. Due to the earth gravity of 9.8 m/s² ("1g"), the measured acceleration includes a constant component that we filter out (similarly to [103,116], we use a 3rd order Butterworth high-pass filter with a 0.1 Hz cutoff frequency).

We examine two motion properties of the measurements: the average absolute deviation of the acceleration, D, and the dominant frequency of motion, f_m . D quantifies the variability in the a(t) value and is a measure of the "amount of motion". We calculate it as $D = \frac{1}{T} \sum_T (a(t) - \overline{a}(t))$, where

¹Although smartphones include accelerometers, we use dedicated sensing units since the phones' accelerometers have a limited range, restricted sampling rate control, and high energy consumption (that hinders day-scale trace collection).



Figure 15: Demonstration of obtaining the power generated by a harvester, P(t), from the recorded acceleration, a(t): (a) a(t) recorded by a person walking, (b) the corresponding harvester proof mass displacement, z(t), and (c) the resulting P(t) for harvester H1.

 $\overline{a}(t)$ denotes the average of a(t) over time interval T. We obtain f_m by determining the maximum spectral component of the Fourier Transform of a(t).

4.1.3 Harvesting Rates and Data Rates

We calculate the power generated by a harvester, P(t), subjected to acceleration a(t), using the following procedure based on the approaches developed in [116, 118]. We first convert a(t) to proof mass displacement, z(t), using the Laplace-domain transfer function:

$$z(t) = \mathcal{L}^{-1}\{z(s)\} = \frac{a(s)}{s^2 + (2\pi \cdot f_r/Q)s + (2\pi \cdot f_r)^2}$$

Next, to account for the Z_L , we limit z(t) using a Simulink limiter block. The power P(t) generated by the harvester is then determined as

$$P(t) = b \cdot \frac{dz(t)^2}{dt}^2.$$



Figure 16: The average power generated by a harvester, \overline{P} , from the same motion (human running) for different combinations of harvester resonant frequencies, f_r , and harvester damping factors, b.

The average of the P(t) is denoted by \overline{P} .

We implemented this procedure in MATLAB and Simulink. Fig. 15 shows an example of obtaining P(t) for a particular a(t). The a(t) values were recorded by a sensing unit carried by a walking person (Fig. 15(a)), and the z(t) and P(t) values were obtained using the procedure described above for the harvester H1.

To characterize the performance of wireless mobile systems, we calculate the data rates, r, that a node would be able to maintain when harvesting the generated \overline{P} . We assume, similar to [116], that the harvester energy conversion efficiency, η , is 20%. As a node's cost to communicate, we use $c_{tx} = 1$ nJ/bit (corresponding to energy spending of ultra-low-power transceivers envisioned for the EnHANTs). Hence, $r = \eta \cdot \overline{P}/c_{tx} = 2 \cdot 10^5 \cdot \overline{P}$ Kb/s.

4.1.4 Optimizing Harvester Parameters

Finding the optimal harvester parameters k and b is difficult, since it requires optimizing over a multi-dimensional surface of unknown geometry [103]. For example, Fig. 16 shows the average power, \overline{P} , values calculated from one set of a(t) measurements (corresponding to 20 seconds of a person running) for different $\{f_r, b\}$ combinations. To determine the optimal harvester parameters for short a(t) samples, we implemented an *exhaustive search* algorithm. This algorithm considers a large number of k and b combinations, obtains the corresponding \overline{P} (using the procedure described in Section 4.1.3), and selects the k and b combination that maximizes \overline{P} .

The exhaustive search algorithm is time-consuming even for relatively short a(t) samples. For optimizing harvester parameters with longer a(t) samples, we implemented a simplified procedure

Table 3: Object motion measurements.

Scenario	\overline{P}
Taking a book off a shelf	$< 10 \ \mu W$
Putting on reading glasses	$<\!10~\mu W$
Reading a book	$<\!10~\mu W$
Writing with a pencil	10–15 μW
Opening a drawer	10–30 μW
Spinning in a swivel chair	${<}10~\mu{\rm W}$
Opening a building door	$<1 \ \mu W$
Shaking an object	$>3{,}000~\mu{\rm W}$

developed in [116]. The procedure first determines the k value that matches the harvester's f_r to the dominant frequency in the a(t) sample, f_m . Specifically, the procedure selects k such that

$$k = \frac{f_r^2 \cdot m}{(2\pi)^2} = \frac{f_m^2 \cdot m}{(2\pi)^2}$$

It then considers a relatively large number of b values and selects the b that maximizes \overline{P} .

4.2 Object Motion Energy

Everyday motions generate large amounts of energy, but not all of that energy can be harvested by an inertial harvester. In this section, we provide some observations regarding the energy availability associated with the *motion of objects*. We conducted extensive measurements, recording a(t) and calculating \overline{P} for a wide range of motions. Our measurements included performing everyday activities with a variety of objects (see Table 3), shipping a FedEx box containing a sensing unit from Houston, TX to New York City, NY, transporting sensing units in carry-on and checked-in airplane luggage, and taking sensing units on cars, subways, and trains.

Below, we present observations based on our measurements. To put the \overline{P} values in perspective, we note that *human walking* typically corresponds to $120 \leq \overline{P} \leq 280 \ \mu\text{W}$, as we will describe in Section 4.3.

• Only periodic motion is energy-rich – Due to the filter properties of inertial harvesters (see Fig. 13(b), for example), a motion needs to be *periodic* to be "harvestable". The vast majority of common object motions are not periodic, and hence the corresponding energy availability is low. For example, we attached a sensing unit to a book and observed that when

the book was being taken off the shelf, read, or put back on the shelf, $\overline{P} < 10 \ \mu\text{W}$. For a sensing unit attached to a pencil used by a student to write homework, $10 \le \overline{P} \le 15 \ \mu\text{W}$. Even high-acceleration non-periodic motions, such as a plane landing and taking off, and an accelerating or decelerating car, correspond to only limited energy availability ($\overline{P} < 5 \ \mu\text{W}$). For example, when a unit was placed in a bag checked in on a 3 hour 13 minutes flight, the recorded a(t) showed that the luggage was subjected to varying high-acceleration motions, but the \overline{P} did not exceed 5 μ W even during the most turbulent intervals of the flight.

- Damped (softened) object motion is energy-poor The motion of many objects in our environment is *damped* for human comfort (e.g., by door dampers, cabinet drawer dampers, and springs in swivel chairs). In such cases, most of the motion energy is absorbed in the dampers and only small amounts can be harvested by external sticker form factor harvesters (such as, for example, EnHANTs). Opening and closing a drawer, spinning a swivel chair, and opening and closing a door of a building corresponded to 10 ≤ P ≤ 30 μW, 1 ≤ P ≤ 6.5 μW, and P < 1 μW, respectively. This suggests that wireless nodes embedded in objects such as doors and drawers should *integrate motion energy harvesters with the mechanical dampers*.
- Purposeful object motion can be extremely energy-rich Periodic shaking of objects can generate a relatively large amount of energy in a short time (as demonstrated by *shake flashlights*). In our experiments, purposeful shaking of a sensing unit corresponded to *P* of up to 3,500 μW, that is, 12–29 times more than the power for walking. In IoT applications with mobile nodes, this can be useful for quickly recharging battery-depleted nodes.

4.3 Human Motion Energy

We now examine a dataset with over 40 participants performing 7 common human motions in unconstrained environments. The dataset was collected in [111] and used for activity recognition, rather than energy characterization. We first introduce the study. Then, we characterize the energy availability for different motions, the variability in motion properties among sensing unit placements and participants, and the dependence of energy availability on the participant's physical parameters.



Figure 17: Characterization of kinetic energy for common human activities, based on a 40-participant study: (a) average absolute deviation of acceleration, D, (b) dominant motion frequency, f_m , and (c) power harvested by an optimized inertial harvester, \overline{P} .

4.3.1 Study Summary

The dataset we examine [111] contains motion samples for 7 common human activities – relaxing, walking, fast walking, running, cycling, going upstairs, and going downstairs – performed by over 40 participants and recorded from the 3 sensing unit placements shown in Fig. 14(b). For each 20 second motion sample in [111], we use the acceleration trace to calculate D, f_m , \overline{P} , and r. To obtain \overline{P} , we use the exhaustive search harvester optimization algorithm described in Section 4.1.4. By determining the best harvester for each motion, we can offer important insights into the harvester design.

To validate the data from [111], we replicated the measurements with our sensing units. The results of our measurements were consistent with the provided data. We note that the f_m values calculated for the different motions in the dataset are consistent with the physiology of human motions.

The statistics of the calculated D, f_m , and \overline{P} are summarized in the boxplots in Fig. 17. For

Activity	Sensing unit	# subjects	Median f_m		Median r		
	placement		(Hz)	25 th per-	Median	75 th per-	(Kb/s)
				centile		centile	
	Trouser pocket	42	N/A	1.0	3.1	4.8	0.6
Relaxing	Waist belt	42	N/A	0.3	2.4	4.8	0.5
	Trouser pocket	42	N/A	0.2	1.4	5.9	0.3
	Shirt pocket	42	1.9	128.6	155.2	186.0	31.0
Walking	Waist belt	42	2.0	151.8	180.3	200.3	36.0
	Trouser pocket	42	2.0	163.4	202.4	274.5	40.4
	Shirt pocket	42	2.8	724.2	813.3	910.0	162.6
Running	Waist belt	41	2.8	623.5	678.3	752.8	135.6
	Trouser pocket	42	2.8	542.3	612.7	727.4	122.5
	Shirt pocket	30	3.5	37.4	52.0	72.3	10.4
Cycling	Waist belt	29	3.8	36.3	45.4	59.2	9.1
	Trouser pocket	30	1.1	35.6	41.3	59.5	8.3

Table 4: Energy budgets and data rates based on measurements of common human activities.

each of the 7 motions the leftmost (black), middle (red), and rightmost (blue) boxes correspond to shirt pocket, waist belt, and trouser pocket sensing unit placements, respectively. For each motion and sensing unit placement, the number of participants that had a(t) samples appears on the top of Fig. 17(a). On each box, the central mark is the median, the edges are the 25th and 75th percentiles, the "whiskers" extend to cover 2.7σ of the data, and the outliers are plotted individually. In Table 4 we separately summarize the motion energy characterizations and the data rates for 4 important motions.

4.3.2 Energy for Different Activities

Relaxing: As expected, almost no energy can be harvested when a person is not moving ($\overline{P} < 5 \ \mu W$).

Walking and fast walking: Walking is the predominant periodic motion in normal human lives and thus particularly important for motion energy harvesting. For walking, the median \overline{P} is 155 μ W for shirt pocket sensing unit placement, 180 μ W for waist belt placement, and 202 μ W for trouser pocket placement. These \overline{P} values are in agreement with previous, smaller-scale, studies of motion energy harvesting for human walking [39,103]. In comparison, *indoor light energy* availability is on the order of 50–100 μ W/cm². Considering harvester energy conversion efficiency estimates [30,116], a similarly sized harvester would harvest more energy from walking than from indoor light. Fast walking (which was identified as "fast" by the participants themselves) has higher D and f_m than walking at a normal pace (Fig. 17) and generates up to twice as much \overline{P} .

Running: Running, an intense repetitive activity, is associated with high D and f_m (Fig. 17(a,b)),

and hence results in $612 \leq \overline{P} \leq 813 \ \mu W$.

Cycling: For the examined sensing unit placements, cycling generates relatively little energy – the median \overline{P} values are 41–52 μ W, 3.7–3.9 times less than the \overline{P} for walking. While the high cadence of cycling motion results in relatively high f_m (Fig. 17(b)), a harvester not on the legs will be subject to only small displacements, resulting in small values of D (Fig. 17(a)) and \overline{P} (Fig. 17(c)). For cycling IoT applications, harvester placements on the lower legs should be considered.

Walking upstairs and downstairs: Our examination demonstrates that human exertion (perceived effort and energy expenditure) does not necessarily correspond to higher motion energy harvesting rates. While people exert themselves more going upstairs, the \overline{P} for going downstairs is substantially higher than for going upstairs. Specifically, for the downstairs motion, the median \overline{P} is 1.78 times higher than the upstairs motion for shirt unit placement, 2.1 times higher for waist placement, and 1.65 times higher for trouser placement. Although counterintuitive, going downstairs is associated with higher magnitudes of motion and higher motion frequencies (Fig. 17(a,b)), which leads to the higher \overline{P} . We observed the disconnect between perceived human effort and energy harvesting rates in our own measurements as well, where we noted that highly strenuous activities, such as push-ups and sit-ups, resulted in higher \overline{P} values than non-strenuous walking at a normal pace.

4.3.3 Consistency of Dominant Motion Frequency

To maximize power output, the resonant frequency of a harvester, f_r , should "match" the dominant frequency of motion, f_m . In this section, we comment on the variability in f_m and provide important observations for harvester design. Due to space constraints, we leave the study of harvester sensitivity to different design parameters to future work.

Consistency among sensing unit placements: The same motion will result in a different f_m depending on the sensing unit's placement on the human body [39, 116]. We observed this in measurements that we conducted, especially for sensing units attached to the lower legs and lower arms. However, for the sensing unit placements examined in this section (i.e., shirt pocket, waist belt, and trouser pocket), the same motion resulted in similar f_m values, as can be seen in Fig. 17(b). These placements are on or near the torso, and are subjected to similar stresses. Cycling is an

exception; the f_m for the trouser pocket placement is different from the other placements. Because the body is in a sitting position, the stresses experienced by the legs and the torso are different, and f_m differs for the different placements.

The uniformity of f_m offers valuable hints for energy harvesting node designers. People are likely to keep many objects that will become energy harvesting communicating Internet of Things nodes (for example, keys, wallets, and cell phones) in pockets located in places that correspond to the placements we examine. This suggests that a harvester will perform well regardless of where a person chooses to carry such an object.

Inter-participant consistency: For common periodic motions, such as walking and running, the f_m values are relatively consistent among the different participants. The 25th and 75th percentiles of the participants' f_m values are separated by only 0.15 Hz for walking and by only 0.3 Hz for running. For less commonly practiced motions (cycling, going upstairs, going downstairs), the values of f_m are less consistent, but are still somewhat similar. This consistency indicates that an all-purpose harvester designed for human walking or running will work well for a large number of different people. The next section examines whether harvesters can be tuned to particular human parameters.

4.3.4 Dependency on Human Height and Weight

We examine the dependency of energy availability on human physiological parameters. We correlate D, f_m , and \overline{P} obtained for different motions and different participants with their height and weight data from [111].² The participants' height range was 155–182 cm, and their weights range was 44–65 kg. We verified that, in agreement with general human physiology studies, the participants' height and weight are strongly positively correlated ($\rho = 0.7$, p < 0.001).

As indicated above, for many activities f_m is consistent among different participants. Yet, we additionally observed f_m dependencies on human physiology. For many of the activities we examined, we determined *negative correlations of* f_m with the participants' height and weight. When walking, running, and going upstairs and downstairs, heavier and taller people took fewer steps per time interval than lighter and shorter people.

 $^{^{2}}$ The dataset [111] is also annotated with participants' age and gender. However, the age range (20 to 23 years) and the number of females (10 participants) are insufficient for obtaining statistically significant correlations.

For example, for going upstairs with waist belt sensing unit placement, f_m and the participant's height are correlated as $\rho = -0.34$ (p < 0.05, n = 39). When going upstairs, the taller half of the participants made, on average, 9 fewer steps per minute (0.15 Hz) than the shorter half ($f_m = 1.85$ and 2.05 Hz, correspondingly). For running, with trouser pocket placement, f_m and the participant's weight are correlated as $\rho = -0.46$ (p < 0.01, n = 39). When running, the heavier half of the participants made, on average, 18 fewer steps per minute (0.3 Hz) than the lighter half. This suggests that future harvester designs may benefit from targeting harvesters with different f_r values for human groups with different physiological parameters. For example, different harvesters may be integrated in clothing of different sizes.

Generally, motion energy availability increases as f_m increases [61]. However, in human motion, other dependencies may additionally come into play. In our study, for running with trouser pocket sensing unit placement, we determined a positive correlation between D and participants' height $(\rho = 0.35, p < 0.05, n = 38)$ and a positive correlation between \overline{P} and participants' height $(\rho = 0.38,$ p < 0.01, n = 38). For the taller half of the participants, the average \overline{P} is 20% higher than for the shorter half (704 and 582 μ W, respectively). Studies with larger number of participants, wider participant demographics, and wider range of participant parameters will most likely identify several additional dependencies. This will allow harvester designers to develop harvesters for different demographics, as well as to provide guarantees on the performance of different harvesters based on different human parameters.

4.4 Long-term Human Mobility

The results presented in the previous section are based on short motion samples from an activity recognition dataset. In this section, we present results of our own, *longer-term, motion measurements*. We describe our set of day-long human mobility measurements and discuss energy budgets and generation process properties.



Figure 18: Motion energy characterization for a 3 hour run: (a) the absolute deviation of acceleration, D, as a function of time, (b) dominant motion frequency, f_m , as a function of time, and (c) the distribution of the corresponding harvested power, P(t).

4.4.1 Prolonged Activities

To study motion energy properties over time, we collected a set of measurements of longer activity durations (over 20 minutes). We considered long walks, bike rides, runs, and other activities, performed in normal environments (i.e., not on a treadmill or a stationary bike). To the best of our knowledge, the properties of motion of longer samples have not been analyzed before.

The measurements demonstrate that for prolonged activities, D, f_m , and P(t) vary substantially over time. This variability is related to physiological parameters, such as changes in cadence or posture over time due to fatigue, and changes in the surrounding environment, such as traffic lights, terrain changes, or pedestrian traffic. For example, Fig. 18 shows D, f_m , and P corresponding to a 3 hour run, calculated for 1-second a(t) intervals. In this trace, the average D changes subtly over time (Fig. 18(a)), and f_m varies continuously in the 2.6–3.4 Hz range (Fig. 18(b)). Correspondingly, while the mean P(t) is 550 μ W, the 10th–90th percentiles of P(t) span the range of 459–710 μ W (Fig. 18(c)).

The variability of P(t) throughout an activity suggests that node energy management policies are essential even for specifically targeted energy harvesting node applications, such as nodes for fitness runners or cyclists. In the following sections, we demonstrate even more variability in P(t)for the regular everyday human mobility patterns.

rouume	<i>.</i> 0.							
Parti-	Occupation and	#	Total	Optimized harvester		r_d ,	\overline{P}^{H4} (μ W),	% ON,
cipant	commute	Days	dur.(h)			avg	min/avg/max	min/avg/max
				\overline{P} (μW),	$\overline{P_d}$ (μ W),	(Kb/s)		
				min/avg/ma	x min/avg/ma	x		
M1	Undergraduate stu-	5	60.4	6.9 / 13.8	4.8 / 6.5 /	1.3	5.0 / 8.5 / 10.9	5.4 / 9.9 / 12.2
	dent, living on cam-			/ 17.3	8.1			
	pus							
M2	Undergraduate stu-	3	27.7	23.3 / 29.0	8.4 / 11.5	2.3	17.1 / 19.6 / 24.5	13.6 / 16.1 / 18.4
	dent, commuting to			/ 38.2	/ 17.7			
	campus			,				
M3	Undergraduate stu-	9	62.0	2.4 / 7.16	0.6 / 2.02	0.4	2.0 / 5.8 / 12.2	3.6 / 6.0 /9.95
	dent, living on cam-			/ 13.4	/ 3.6			
	pus			,				
M4	Graduate student,	7	80.1	1.4 / 11.98	0.6 / 5.6 /	1.1	1.4 / 11.98 / 25.3	2.8 / 12.7 / 18.1
	commuting to cam-			/ 25.3	10.7			
	pus							
M5	Software developer,	1	11.0	16.3	7.5	1.5	15.9	11.5
	commuting to office							

Table 5: Energy budgets, variability, and data rates based on collected traces for daily human routines.



Figure 19: Kinetic energy for normal daily human routine: (a) acceleration, a(t), recorded over 11 hours for participant M5, and (b) the power harvested, P(t).

4.4.2 Day-Long Human Mobility

To determine the daily energy available to a mobile node with an inertial harvester, we collected acceleration traces from different participants during their normal daily routines. Our results are based on over 200 hours of acceleration information we obtained for 5 participants for a total of 25 days. Some information about the participants is provided in Table 5. The participants were instructed to carry a sensing unit in any convenient way. Thus, the measurements correspond to the motion that a participant's keys, a mobile phone, or a wallet would experience.

Fig. 19 shows the a(t) for a day-long trace of participant M5, and the corresponding P(t). For all the collected traces, the dominant motion frequency, f_m , range is 1.92–2.8 Hz, corresponding to human walking.

The calculated energy budgets are summarized in Table 5. We calculated \overline{P} , the average power a harvester would generate over the length of the trace, as well as $\overline{P_d}$, the average power a harvester would generate over a 24-hour interval. To calculate $\overline{P_d}$ we assumed that when the sensing unit did not record data (e.g., before the participants got dressed for school or work), it was stationary and that a harvester would not generate energy during these intervals. Specifically, for a T hour-long measurement trace, $\overline{P_d} = \overline{P} \cdot T/24$. For each of the participants, Table 5 summarizes the minimum, average, and maximum \overline{P} and $\overline{P_d}$ over the different measurement days, and the data rate r_d that a node would be able to maintain continuously throughout a day when powered by the harvested \overline{P}_d . For completeness, for all participants we additionally calculate \overline{P}^{H4} , the average power a particular harvester, same for all participants (in this case, the harvester calculated based on the traces for the participant M4), would harvest.

For most participants, an inertial harvester can provide sufficient power to continuously maintain a data rate of at least 1 Kb/s (i.e., $\overline{P_d} > 5 \ \mu W$). This is comparable with the data rates estimated in Section 3.3 for nodes with a similarly sized *light* harvester in indoor environments (not exposed to outdoor light).

The majority of inter-participant and inter-day differences seem to relate to the participants' amount of walking. For example, the participant M2, whose \overline{P} and $\overline{P_d}$ values are higher than the values for the other participants, has a relatively long walk to the office, and walks frequently between two different offices in the same building. Other factors (e.g., unit placement, amount of daily activity as perceived by the participants) appear to be only of secondary importance. We note that the majority of traces that correspond to $\overline{P_d} < 5 \ \mu$ W (and thus $r_d < 1 \ \text{Kb/s}$) correspond to participants working from home. Overall, daily routines that involve a lot of walking correspond to relatively high levels of energy availability.

4.4.3 Harvesting Process Variability and Properties

The amount of energy that can be harvested varies widely throughout the day. As shown in Section 4.3, *walking* generates substantial amounts of energy, while being stationary generates little.



Figure 20: Motion energy harvesting process variability for participant M1: (a) the percentage of the total energy harvested at each of the demonstrated power levels, P(t), and (b) the *percentage* of time the power is harvested at each of the demonstrated P(t) (notice that for $0 \le P(t) \le 15$, the value is 91%).

Physiological studies (e.g., [66]) have shown that people are at rest a majority of the time. Correspondingly, in our measurements, P(t) is low for most of the day and over 95% of the total energy is collected during only 4–7% of a day. For example, Fig. 20 shows, for participant M1, the percentage of the total energy that would be harvested over different ranges of P(t) and the *percentage of the time* that the harvester would generate these P(t) values. For this participant, the harvester would generate $P(t) < 15 \ \mu\text{W} \ 91\%$ of the time, but only 6.1% of the total energy would be harvested during this time.

Consider an ON/OFF representation of the energy harvesting process, $P_{\text{onoff}}(t)$, where $P_{\text{onoff}}(t) \leftarrow 1$ ("ON") if P(t) > THR, and $P_{\text{onoff}}(t) \leftarrow 0$ ("OFF") otherwise. For the analysis below, we empirically set $THR = 10 \ \mu\text{W}$. The results are similar for $10 \leq THR \leq 40 \ \mu\text{W}$. For all participants, it can be seen that the process is ON for less than 20% of the time (Table 5). Note that the participants do not lead sedentary lifestyles; their activity patterns are well in line with general health guidelines. However, the recommend 30 minutes of physical activity per day correspond to only 9% of an 11 hour measurement trace. Additionally, the typical *duration* of ON intervals is short – on the order of seconds. While some of the ON intervals are long (over 200 seconds), the vast majority of the ON intervals (78.5–89.0%) are shorter than 30 seconds; the median ON intervals are 5–9.5 seconds. The longer ON intervals correspond to commuting (e.g., walking from a subway station to a campus building), and represent only 1–3% of the ON intervals.

In summary, P(t) is low for the majority of the time, and when it does become high, it stays high for only a brief period of time. This emphasizes the need for energy management policies for communicating wireless nodes powered by this energy.

4.5 Conclusions and Future Work

In this chapter, we considered kinetic (motion) energy for ultra-low-power energy harvesting nodes. Based on our measurements, we provide observations regarding the energy of *object motion*. We also thoroughly study the energy associated with *human motion*. To characterize human motion energy availability, we use the results of our measurement campaign that include 200 hours of acceleration traces from day-long human activities. Moreover, we use a dataset of 7 common human motions performed by over 40 participants [111]. With regards to object motion, our study demonstrates relatively low energy availability levels for most object motions. In particular, the study highlights the need for periodic motion which is *common in human motion but is uncommon for many objects*. For human motion, we demonstrate the importance of different human parameters. For example, we show that the taller half of the participants can harvest on average 20% more power than the shorter half. Additionally, we show that the energy availability from normal human routines is compatible to energy availability from indoor lights in enclosed environments. Moreover, we demonstrate that the power generation process associated with human motion is highly variable, with only brief intervals of high power levels.

As part of future work, we plan to make publicly available (e.g., publish via CRAWDAD) the developed modeling code and the motion dataset we assembled. To the best of our knowledge, this will be the first publicly available long-term human motion acceleration dataset. Additional future work may include *joint measurements* of motion and light energy availability. Such studies would be particularly important for designing energy harvesting nodes relying on emerging *combined* (*multi-source*) energy harvesters, such as [4].

Chapter 5

Resource Allocation Algorithms for Energy Harvesting Nodes

In this chapter, we formulate and examine resource allocation problems for energy harvesting nodes. As demonstrated in Chapters 3 and 4, the energy available to energy harvesting nodes varies in time. Inspired by the needs of tracking and monitoring applications (such as some of the applications proposed for EnHANTs), to avoid node outages and to gain control over node and network behavior, we aim to allocate energy harvesting nodes' resources in a uniform way with respect to time. That is, we aim to obtain "smooth" resource allocations which ensure that the nodes and the networks maintain a certain level of performance at all times despite the variations in the energy availability.¹ To formulate the resource allocation problems, we use the utility maximization and the lexicographic maximization frameworks. These frameworks are typically applied to achieve fairness among nodes [7,19,48,55,67]. We apply them to achieve time-fair resource allocations.

We mainly consider *deterministic energy profile* and *stochastic* environmental energy models (see Section 1.2), and focus on *single node* and *link* scenarios. For the *deterministic energy profile* model, we formulate optimization problems and introduce algorithms for solving the formulated problems. We used the algorithms to obtain numerical results for various cases, using, as energy inputs, light

¹The need for policies that enable such "smooth" behavior in energy harvesting devices has been noted by many researchers [19, 46, 64, 102].

energy traces we presented in Chapter 3. We also examine a range of simple policies, for which we provide performance guarantees. For the *stochastic* energy model, we consider the case in which the energy inputs are i.i.d. random variables. We formulate resource allocation problems as *average cost Markov Decision Process* (MDP) problems. We study the *uniform discretization of the problem*, for which we obtain a bound on the performance degradation due to discretization. We introduce algorithms for solving the problem, and provide performance guarantees for a set of simple policies.

Additionally, we briefly consider *model-free* approaches, for which we examine the performance of a set of simple online policies with the kinetic energy traces we presented in Chapter 4.

Several contributions of this chapter are based on joint work with A. Wallwater or A. Bernstein. A. Wallwater contributed to problem formulations and analysis of some of the algorithms for the deterministic profile energy model. A. Bernstein contributed to problem formulations and analysis of some of the algorithms for the stochastic energy model.

In this chapter, we first present the model and the frameworks (Section 5.1). Then, we present the formulated problems and the solution algorithms for the deterministic profile energy model (Section 5.2) and the stochastic energy model (Section 5.3). Finally, we briefly examine the performance of online policies with kinetic energy traces (Section 5.4).

We previously presented some of the results that appear in this chapter in [21, 30, 31].

5.1 Model

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) ($0 \le B(i) \le C$). We denote the initial and the final energy levels by B_0 and B_K , respectively. Most policies we consider aim to ensure *energy neutrality* – full spending, yet not over-spending, of the environmental energy, i.e., $B_K = B_0$ [19,46]. For most evaluations, we use $B_0 = B_K = C/2$.

Table 6 summarizes the notation used in this chapter.

We focus on the *deterministic profile* and *stochastic* environmental energy models, introduced in Section 1.2. We let D(i) denote the *environmental energy exposure rate* for a node in a time slot

	Table 6: Nomenclature.
i, K	Time slot index, and number of time slots
C	Energy storage capacity (J)
$B(i), B_0, B_K$	Energy storage state, initial, and final levels (J)
D(i)	Environmental energy exposure rate (W)
Q(i)	Energy harvesting rate (W)
8	Energy spending rate (J/slot)
\hat{Q}	Total energy to be allocated (J)
Δ	Quantization resolution (J)
r	Data rate (bits/s)
$c_{\mathrm{tx}}, c_{\mathrm{rx}}$	Energetic costs to transmit and to receive (J/bit)
$U(\cdot)$	Utility function
Ζ	Objective function value
T, T_L	Node downtime, link downtime

i. The energy profile is a vector $\{D(0), D(1), ..., D(K-1)\}$ that corresponds to the environmental energy exposure rates at different time slots. For the stochastic energy models, the environmental energy exposure rate is a stochastic process $\{D(i)\}$. For simplicity, we assume that $\{D(i)\}$ are independent identically distributed (i.i.d.) random variables², and use D to denote the representative variable for D(i), with p_D being its probability density function (pdf).

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 non-linear function of D(i). 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, or general energy storage. Functions q(D(i), B(i)) for a capacitor, derived from capacitors' electric properties, were previously shown in Fig. 3. To derive numerical results for nonlinear energy storage, we use $q(D(i), B(i)) = D(i) - D(i) \cdot (B(i) - C/2)^2/(\beta_{nonlin} \cdot (C/2)^2)$, where β_{nonlin} is the energy storage nonlinearity parameter.³ These functions have properties similar to the capacitor-based energy harvesting system's q(D(i), B(i))functions previously shown in Fig. 3.

The node energy spending rate is denoted by s(i). The energy storage evolution of an energy harvesting device can be expressed as:

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

²In [22], consider D(i) with other properties.

³We note that Q(i) = D(i) for $\beta_{nonlin} \to \infty$.

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 endpoints' 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 assume that the link endpoints are 'saturated', that is, always have information to send to each other. We denote the costs to transmit and receive a bit by c_{tx} and c_{rx} . The constraints relating energy spending rates and data rates for a link (u,v)at time slot *i* are:

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

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 [7, 19, 48, 55, 67]. In this work 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 the objective function, $Z, Z = \sum_{i=0}^{K-1} U(s(i))$ (for a single node) or Z = $\sum_{i=0}^{K-1} [U(r_u(i)) + U(r_v(i))]$ (for a link), where $U(\cdot)$ are concave non-decreasing twice-differentiable continuous functions (e.g., $U(\cdot) = \log(\cdot), U(\cdot) = \sqrt{(\cdot)}, U(\cdot) = (\cdot)^{1-\alpha}/(1-\alpha), \alpha > 1)$.⁴ In general, the solutions obtained by applying the two frameworks are not the same. The solutions are identical in certain cases, such as those we examine in Lemma 1 and in Observation 3.

We evaluate the energy harvesting adaptive policies in terms of the objective function values, Z. To derive numerical results, we use $U(\cdot) = \log(\cdot)$ or $U(\cdot) = \log(1 + (\cdot))$, where the " (\cdot) " is s(i) in node

⁴We 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 $\alpha \to \infty$ [62].

scenarios, and r(i) in link scenarios. We use superscripts to indicate the policy under which Z was obtained (e.g., Z^{opt} for the optimal policy, Z^{cr} for the constant rate policy). We additionally consider the *downtimes* of nodes and links, namely, the fraction of slots the node or the link do not have energy expenditure allocated. We denote the downtime of a node u by T_u , where $T_u = |\{i|s_u(i) = 0\}|/K$, and the downtime of a link (u, v) by $T_{L(u,v)}$, where $T_{L(u,v)} = |\{i|r_u(i) = 0, r_v(i) = 0\}|/K$. Additional performance metrics used with the model-free approaches in Section 5.4 are introduced in the section itself.

5.2 Deterministic Energy Profile

In this section, we consider the *deterministic profile* energy model, which is similar to the models studied in [19, 46, 65]. We formulate optimization problems that apply to both *linear* and *nonlinear* energy storage for a *single node* (Section 5.2.1) and for a *link* (Section 5.2.2). We introduce algorithms of different complexity for different settings, and additionally consider a set of simple algorithms, for some of which we obtain performance guarantees.

The proofs for Section 5.2.1 appear in Appendix 5.A.1. The proofs for Section 5.2.2 appear in Appendix 5.A.2.

5.2.1 Single Node: Optimizing Node 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:

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

s.t.:
$$s(i) \le B(i) \ \forall \ i$$
 (4)

$$B(i) \le B(i-1) + Q(i-1) - s(i-1) \ \forall \ i \ge 1$$
(5)

$$B(i) \le C \ \forall \ i \tag{6}$$

$$B(0) = B_0; \ B(K) \ge B_K \tag{7}$$

$$B(i), s(i) \ge 0 \ \forall \ i. \tag{8}$$

Recall that U(s(i)) is a concave non-decreasing function. Recall, additionally, that for *linear energy* storage, Q(i) = D(i), and for *nonlinear energy storage*, Q(i) = q(D(i), B(i)). Constraint (4) ensures that a node does not spend more energy than it has stored, (5) and (6) represent the energy storage evolution dynamics, and (7) sets the initial and final energy storage levels to B_0 and B_K .

Time Fair Lexicographic Assignment (TFLA) Problem:

Lexicographically maximize:
$$\{s(0), ..., s(K-1)\}$$
 (9)
s.t.: constraints $(4) - (8)$.

Fig. 21 shows an example of node energy allocation vectors $\{s(i)\}$ obtained by solving the TFU and the TFLA problems. Fig. 21(a) 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 1). Fig. 21(b) 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.⁵

The optimal solution to the TFU Problem is bounded as follows:

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

For linear energy storage (q(D(i), B(i)) = D(i)), i.e., a battery), we refer to the TFU and the

⁵The solutions were obtained for the following parameters: $C = 0.5 \cdot \hat{Q}$, $B_0 = B_K = 0.4 \cdot C$, $U(s(i)) = \log(s(i))$, $\beta_{\text{nonlin}} = 1.05$.



Figure 21: (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).

TFLA problems as **TFU-LIN** and **TFLA-LIN** problems. For these problems, we obtain the following Lemma.

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

Solution Algorithms

For solving the formulated problems, 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).

Algorithm 1 Time Fair Rate Assignment (TFR).	
$h(i, B) \leftarrow -\infty, s(i) \leftarrow 0 \forall i < K, \forall B;$	
$h(K,B) \leftarrow -\infty \forall B < B_K;$	
$h(K,B) \leftarrow 0 \ \forall \ B \ge B_K;$	
for $i = K - 1; i \ge 0; i;$ do	
for $B = 0; B \leq C; B \leftarrow B + \Delta;$ do	
for $s = 0; s \leq B; s \leftarrow s + \Delta;$ do	
$\hat{s} \leftarrow s; \ \hat{h} \leftarrow U(\hat{s}) + h(i+1, \min(B+q(D(i), B) - \hat{s}, C);$	
${f if}\hat{h}>h(i,B){f then}$	
$h(i,B) \leftarrow \hat{h}; s(i) \leftarrow \hat{s};$	
return $h(0, B_0)$, and associated $s(i) \forall i$	

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

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

Going backwards from i = K-1, we thus obtain a vector $\{s(0), ..., s(K-1)\}$ 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 $\{i, B(i)\}$. Maximizing an instance of h(i, B(i)) requires considering all s(i) such that $s(i) \leq B(i) \leq C$. 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 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 [7]. 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}/Δ 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}/Δ , 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).

$A_{fix} \leftarrow \varnothing; \ s(i) \leftarrow 0 \ \forall \ i;$
while $\overline{A_{fix}} \neq \emptyset$ do
for $i = 0; i \le K - 1; i + +;$ do
$\mathbf{if} \ i \in \overline{A_{fix}} \ \mathbf{then}$
$\tilde{s}(j) \leftarrow s(j) \ \forall \ j \in [0, K-1]; \ \tilde{s}(i) \leftarrow \tilde{s}(i) + \Delta;$
$valid \leftarrow \mathbf{check_validity}(\tilde{s});$
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 $i = 1; i \le K; i + +;$ 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

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.⁶ 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 $C - B_0 \ge [\sum_i Q(i)] \cdot (1 - 1/K)$.

Lemma 2 When the LS Conditions or the LS-gen Conditions hold, the optimal solution to the TFLA-LIN problem 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 problem, for U(s(i)) that are twice differential strictly concave on $(0, \hat{Q}]$, and that satisfy

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

Examples of $U(\cdot)$ that satisfy (I), (II), (III) include: (I): $U(\cdot) = (\cdot)^{1-\alpha}/(1-\alpha)$ for $0 < \alpha < 1$ [62], (II): $U(\cdot) = \log(\alpha + (\cdot))$ for $\alpha > 0$, used, for $\alpha = 1$, in e.g., [10,21], (III): $U(\cdot) = \log(\cdot)$, used in e.g., [55]. Verifying that the LS Conditions (or the LS-gen Conditions) hold and determining the corresponding optimal policy is computationally inexpensive.

Simple Policies

In addition to solving the formulated problems optimally, we can also use simple policies. We consider *Spend-What-You-Get* (SG) policies, where the node aims to spend all the energy harvested in a slot, that is $s(i) \leftarrow Q(i) \forall i$. Similar policies were proposed in [55]. The SG policies have very

⁶To determine if the LS Conditions hold, a node needs to know $\{Q(0), ..., Q(K-1)\}$, 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 (see Chapter 3) and do not calculate their energy profiles.
low complexity. For the deterministic profile energy model, we additionally consider Constant Rate (CR) policies, where a node spends energy at the same rate in all time slots, that is, $s(i) \leftarrow s_{cr} \forall i$. Similar policies were proposed in [19].

We now provide approximation ratios for the CR and SG policies (Propositions 1 and 2, correspondingly). Proposition 2 applies to the general energy storage model, while Proposition 1 is restricted to the linear energy storage model.

Proposition 1 Under the CR policy, for $B_K = B_0$ (energy neutrality), $B_K \leq \sum_{i=0}^{K-1} Q(i)$, and $U(s(i)) = \log(1+s(i))$,

$$Z^{\rm cr} \ge Z^{\rm opt} \cdot \left(B_0 / \sum_{i=0}^{K-1} Q(i) \right). \tag{10}$$

For example, it follows that for $B_0 = C/2$ and $\sum_{i=0}^{K-1} Q(i) = 3 \cdot C/4$, the CR policy is a 1.5-approximation algorithm.

Proposition 2 Under the SG policy, for $U(s(i)) = \log(M + s(i))$ where M is a constant⁷,

$$Z^{\rm sg} \ge Z^{\rm opt} \cdot \log(G\{Q(i) + M\}) / \log(\overline{\{Q(i) + M\}}), \tag{11}$$

where $\overline{\{\cdot\}}$ and $G\{\cdot\}$ denote, respectively, the arithmetic and the geometric means of a sequence.

For example, consider a $\{Q(i)\}$ where in L (out of K) slots Q(i) = A (non-zero constant), and Q(i) = 0 in other slots. Such Q(i) may correspond to the case where the indoor lights are on for a portion of the day. Using Proposition 2, we can show that for $B_K = B_0$ and $U(s(i)) = \log(1 + s(i))$, the SG policy is a K/L-approximation algorithm.

5.2.2 Link: Optimizing Data Rates

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

⁷Note that this proposition can be applied to different utility functions, e.g., to $U(s(i)) = \log(0 + s(i))$ and to $U(s(i)) = \log(1 + s(i))$.

Link Time Fair Utility Maximization (LTFU) Problem:

$$\max_{r_u(i), r_v(i)} \sum_{i=0}^{K-1} [U(r_u(i)) + U(r_v(i))]$$
(12)

s.t.:
$$c_{tx}r_u(i) + c_{rx}r_v(i) \le s_u(i) \forall i$$
 (13)

$$c_{\rm tx}r_v(i) + c_{\rm rx}r_u(i) \le s_v(i) \ \forall \ i \tag{14}$$

u, v: constraints (4) - (8).

Link Time Fair Lexicographic Assignment (LTFL) Problem:

Lexicographically maximize:

$$\{r_u(0), ..., r_u(K-1), r_v(0), ..., r_v(K-1)\}$$
s.t.: (13), (14); u, v : constraints (4) - (8). (15)

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:

Lexicographically maximize: $\{r(0), ..., r(K-1)\}$ (16)

s.t.:
$$r(i) \cdot (c_{tx} + c_{rx}) \le \min(s_u(i), s_v(i)) \forall i$$
 (17)

u, v: constraints (4) - (8)

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

Examples of solutions to the LTFU and LTFL problems are shown in Fig. 22. Fig. 22(a) 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 1) on the same day. Fig. 22(b) shows the data rate allocation vectors $\{r_u(i)\}$ and $\{r_v(i)\}$ obtained by solving the LTFU and the LTFL problems.⁸

⁸The solutions were obtained for the following parameters: $C_u = C_v = 0.5 \cdot \sum_i Q_u(i)$, $B_{0,u} = B_{0,v} = B_{K,u} = B_{K,v} = 0.25 \cdot C_u$, $c_{tx} = 0.1 \text{ nJ/bit}$, $c_{rx} = 1 \text{ nJ/bit}$, $U(r(i)) = \log(r(i))$, and Q(i) = D(i) (linear energy storage).



Figure 22: (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.

We now demonstrate an upper bound on the optimal solution to the LTFU problem (Observation 2). The observation is restricted to the linear energy storage model.

Observation 2 Let \tilde{Z} denote the solution to

$$\max_{r_u(i), r_v(i)} \left\{ U(r_u(i)) + U(r_v(i)) \right\}.$$
(18)

under constraints (13), (14), for $s_u(i) \leftarrow [B_{0,u} - B_{K,u} + \sum_j Q_u(j)]/K$ and $s_v(i) \leftarrow [B_{0,v} - B_{K,v} + \sum_j Q_v(j)]/K$.⁹ The optimal solution is bounded as:

$$Z^{\text{opt}} \le K \cdot \tilde{Z}. \tag{19}$$

In general, the solutions to the LTFL and LTFU problems are not the same. The following Observation identifies a case where the solutions are identical.

Observation 3 When $c_{tx} = c_{rx}$, 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 the **LTFR** algorithm. Over all $\{r_u(i), r_v(i)\}$ such that $c_{tx}r_u(i) + c_{rx}r_v(i) = s_u(i) \leq B_u(i), c_{tx}r_v(i) + c_{rx}r_u(i) = s_v(i) \leq B_v(i)$, the LTFR algorithm determines, for each $\{i, B_u(i), B_v(i)\}$,

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

⁹For example, for $c_{tx} = c_{rx}$, $\tilde{Z} = 2 \cdot U([1/(c_{tx} + c_{rx})] \cdot \min\{[B_{0,u} - B_{K,u} + \sum_{i} Q_u(i)]/K, [B_{0,v} - B_{K,v} + \sum_{i} Q_v(i)]/K\}).$



Figure 23: Comparison of the optimal (OPT) link policy determination (a), and the Decoupled Rate Control (DRC) algorithms (b).

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* 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 $O(K \cdot [K + (\hat{Q}_u + \hat{Q}_v)/\Delta])$.

Decoupled Rate Control (DRC) Algorithms

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 (13) and (14), the nodes obtain a solution to

$$\max_{r_u(i), r_v(i)} U(r_u(i)) + U(r_v(i))$$
(20)

if the LTFU problem is being solved (LTFU-DRC algorithm), and to

$$\max r(i) \tag{21}$$

if the LTFL problem is being solved (**LTFL-DRC** algorithm). These subproblems (each considers a single slot *i*) can be easily solved. 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 $c_{tx} = \rho \cdot c_{rx}$, $\rho > 1$, for the case of $s_v(i) = \gamma s_u(i)$, $0 \le \gamma \le 1$, the optimal solution is either $\{r_u(i), r_v(i)\} = (s_u(i)/(c_{rx} \cdot (\rho^2 - 1)))\{\rho - \gamma, \gamma \cdot \rho - 1\}$ or $\{r_u(i), r_v(i)\} = \{s_v(i)/(2 \cdot c_{rx}), s_v(i)/(2 \cdot c_{tx})\}$. For the LTFL-DRC algorithm, due to (17), the subproblem solution is $r(i) = \min(s_u(i), s_v(i))/(c_{tx} + c_{rx})$.

Fig. 23 demonstrates the difference between solving link problems optimally and applying the DRC algorithms.

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.

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 $R = \max{\{\hat{Q}_u, \hat{Q}_v\}} / \min{\{c_{tx}, c_{rx}\}}$, and that satisfy

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

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

For linear energy storage, even when the LS conditions do not hold, the LTFL-DRC is optimal under the following criteria.

Proposition 3 The LTFL-DRC policy solves the LTFL problem optimally, if for all slots *i*, for $s_u(i)$ and $s_v(i)$ obtained under the OPT policies, $s_u(i) \leq s_v(i)$.

This implies that the LTFL-DRC policy obtains the optimal solution to the LTFL problem when the nodes u and v have the same energy parameters $Q(i), C, B_0$, and B_K .



Figure 24: (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.

In the above-examined LTFU-DRC and LTFL-DRC policies, the nodes' spending rates $s_u(i)$, $s_v(i)$ are determined according to the OPT node policy. We refer to LTFU-DRC as the *node-optimal DRC*, DRC-NOPT. Additionally, we examine several other variants of the DRC: the DRC-SG policy and the DRC-CR policy, where the nodes' spending rates are determined according to the SG or CR policies, correspondingly.

The following observation discusses the downtime under the DRC-SG policy. It applies to the general energy storage model.

Observation 4 Under the DRC-SG policy, $\max[T_u, T_v] \leq T_{L(u,v)} \leq T_u + T_v$.

For example, consider $Q_u(i)$ and $Q_v(i)$ with L non-zero entries. For (u, v) with $Q_v(i) = Q_u(i) \forall i$, $T_{u,v}^L = (K - L)/K$; while for (u, v) with $Q_v(i)$ shifted with respect to $Q_u(i)$, $T_{u,v}^L$ can be as high as $2 \cdot (K - L)/K$.

5.2.3 Numerical Results

This section provides numerical results demonstrating the use of the algorithms described in Sections 5.2.1 and 5.2.2. Measurement traces described in Chapter 3 are used as inputs to the algorithms. In a later Section 6.2, we additionally provide performance evaluation results for some of these policies that we obtained using the EnHANTs testbed.

Fig. 24 shows the optimal energy spending allocation vectors $\{s(i)\}$ for the TFU and the TFLA problems, for different values of energy storage capacity C and initial energy storage state B_0 .¹⁰ The energy profile $\{D(i)\}$ used as an input to these algorithms is shown in Fig. 24(a). It corresponds to

¹⁰The solutions were obtained for the following parameters: $B_K = B_0$ and U(s(i)) = log(s(i)).



Figure 25: Communication rates $r_u(i)$ and $r_v(i)$: (a) obtained by solving the LTFU problem with *nonlinear storage*, and (b) obtained by the LTFU-DRC algorithms for linear energy storage.

the average daily energy profile for the indoor location L-3 (see Table 1). Fig. 24(b) 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 5.2.1, 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. 24(c) 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. 25 shows the numerical results for the *link* data rate determination problems presented in Section 5.2.2. The energy profiles of indoor setups L-1 and L-2 (see Fig. 22(a)) 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. 22. Fig. 25(a) shows the optimal solution to the LTFU problem for *nonlinear energy storage* (for $\beta_{\text{nonlin}} = 1.1$) obtained using the LTFR algorithm. Fig. 25(b) 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.

5.3 Stochastic Energy Models

In this section, we study models in which the amount of energy available in a slot *i* is an *i.i.d.* random process $\{D(i)\}$. *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.

Similarly to the examinations we presented above for the deterministic profile energy model, we are seeking to allocate the energy in a *uniform way with respect to time*. In Section 5.3.1 we formulate the energy allocation problem as an *average cost Markov Decision Process (MDP)*, and propose a way of computing an approximation to an optimal policy using *uniform discretization* of the problem. Our results hold for *linear* and *nonlinear* energy storage models. For the discretized problem, we demonstrate how to calculate discretized policies for *node* and *link* scenarios (Section 5.3.2). We examine a set of simple policies in Section 5.3.3, and provide some numerical results in Section 5.3.4.

The proofs are given in Appendix 5.B.

5.3.1 Optimal Policies and Discretization Bounds

We first formulate the problem of allocating node energy spending as an average cost MDP. The goal is to find an optimal policy, which maximizes the expected average utility. We denote the representative variable for D(i) by D and its probability density function (pdf) by p_D . We denote the state and action spaces of the MDP by $\mathcal{B} = [0, C]$ and $\mathcal{S} = [0, C]$, respectively. For any $b \in \mathcal{B}$ and $s \in \mathcal{S}$, the state transition density is denoted by $p(\cdot|b, s)$. It determines the next energy storage level B(i + 1) given that the current energy storage level is B(i) = b and the spending rate is s(i) = s. This transition density is determined by the energy distribution p_D , the function $q(\cdot, \cdot)$, and (1). A policy π is a collection of decision rules $\pi_i : \mathcal{B}^i \times \mathcal{S}^{i-1} \to \Theta(\mathcal{S})$ which at each time *i* prescribe a probability distribution over the actions ($\Theta(\mathcal{S})$ denotes the probability simplex over the set \mathcal{S}). In particular, we let¹¹ $\lambda_{\pi}(b) \triangleq \lim_{K \to \infty} \mathbb{E}_{\pi}(Z/K) = \lim_{K \to \infty} \mathbb{E}_{\pi}\left(\sum_{i=0}^{K-1} U(s(i))/K\right)$ denote the

 $^{{}^{11}\}mathbb{E}_{\pi}$ denotes the expectation with respect to the probability law induced by the MDP while using policy π , and $\{s(i)\}$ are the spending rates under this policy.

asymptotic expected average utility obtained by starting from state $B_0 = b$ and using a given policy π . The optimal expected average utility is then $\lambda^*(b) \triangleq \sup_{\pi} \lambda_{\pi}(b)$. It is well known (e.g., [73]) that under certain *ergodicity* (or *mixing*) conditions, $\lambda^*(b)$ does not depend on b. In our case, we show that these conditions are satisfied when we impose the following assumption on the energy distribution. Let $Q_b \triangleq q(D, b)$ denote the random variable that represents the harvested energy when the current storage level is b. Consequently, we let p_{Q_b} denote the corresponding density function.

Assumption 1 There exists a finite constant Q_{\max} such that $Q_b \in [0, Q_{\max}]$ for all b. In addition, $p_{Q_b}(y) > 0$ for all b and $y \in [0, Q_{\max}]$.

We note that Assumption 1 will hold for any practical functions q and any random variable D which has a positive density function with finite support.

Under the mixing condition, an optimal policy is a deterministic Markov stationary policy π^* : $\mathcal{B} \to \mathcal{S}$ and can be found by solving the optimality equation $\lambda + J(b) = \mathcal{T}J(b), \ b \in \mathcal{B}$, where \mathcal{T} is Bellman's operator, defined for any bounded function J as

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

The solution (λ^*, J^*) of the optimality equation is such that $\lambda^*(b) \equiv \lambda^*$ and an optimal policy is given by $\pi^*(b) = \operatorname{argmax}_{s \in S} \{U(s) + \int_{\mathcal{B}} p(b'|b, s) J^*(b') db'\}$. However, since our state and action spaces are infinite, there is no practical algorithm to solve the optimality equation.

To address this, we discretize the state and action spaces uniformly, using a fixed discretization parameter Δ . We note that for practical energy-harvesting nodes, uniform discretization is a realistic assumption (e.g., the digital logic that monitors battery levels in the EnHANTs prototypes has a particular resolution Δ). We denote the obtained finite spaces by \mathcal{B}_{Δ} and \mathcal{S}_{Δ} . In particular, if $b \in \mathcal{B}_{\Delta}$, it is a multiple of Δ , and similarly for \mathcal{S}_{Δ} . For the discretized model, we let λ^*_{Δ} and π^*_{Δ} denote the solution of the optimality equation and the corresponding optimal policy, respectively.

We now bound the distance between λ_{Δ}^* in the discretized model and the true optimal expected average utility λ^* (Theorem 1). This bound is an application of the results in [12]. It requires



Figure 26: 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.

an additional Assumption 2 shown below. We note that Assumption 2 holds for any representative random variables D with Lipschitz continuous density function and any Lipschitz continuous function q(D, B).

Assumption 2 The family $\{p_{Q_b}(y)\}$ is Lipschitz continuous in both b and y. In particular, there exists a finite constant $\beta_Q > 0$ such that $|p_{Q_b}(y) - p_{Q_{b'}}(y')| \leq \beta_Q \max(|y - y'|, |b - b'|)$ for all $0 \leq y, y' \leq Q_{\max}$ and $0 \leq b, b' \leq C$.

Theorem 1 Under Assumptions 1 and 2, there exists $\overline{\Delta} > 0$ and β_{λ} such that for all $\Delta \in (0, \overline{\Delta}]$, it holds that $|\lambda^* - \lambda^*_{\Delta}| \leq \beta_{\lambda} \Delta$.

The policy π_{Δ}^* may be computed *offline*. Therefore, the actual choice of the spending rate by a node can be done by using the precomputed function $\pi_{\Delta}^* : \mathcal{B}_{\Delta} \to \mathcal{S}_{\Delta}$.

Remark. The MDP formulation presented above can be extended to a *link* (u, v) by considering the energy harvested in slot *i* by both nodes $D(i) \triangleq (D_u(i), D_v(i))$.

5.3.2 Solution Algorithms

For the discretized problem, we can apply traditional problem solving techniques to obtain the energy allocation policies. We assume that D takes one of M discrete values $[d_1, ..., d_M]$ with probability $[p_1, ..., p_M]$.

For the *node* scenario, the problem can then be formulated as follows.

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



Figure 27: Optimal communication rates $r_u(i)$ (a) and $r_v(i)$ (b), corresponding to different energy storage states, obtained by solving the LSPD problem.

$$\max_{s(i)} \lim_{K \to \infty} \frac{1}{K} \sum_{i=0}^{K-1} U(s(i)).$$
(22)

This discrete time stochastic control process 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

$$h(i, B(i)) = \max_{s(i) \le B(i)D} \mathbb{E}[U(s(i)) + h(i+1, \min[B(i) + q(D(i), B(i)) - s(i), C])]$$
(23)
$$= \max_{s(i) \le B(i)} [U(s(i)) + \sum_{j=1}^{M} p_{d_j} \cdot h(i+1, \min[B(i) + q(d_j, B(i)) - s(i), C])].$$

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 [37]. Although such policy calculations are computationally expensive (the running time of this algorithm is $O([C/\Delta]^2 \cdot M \cdot K))$, a policy needs to be computed only once for a particular distribution of D. Fig. 26 shows 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. 10) is used as the random variable D.¹²

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

Link Spending Policy Determination (LSPD) Problem:

$$\max_{r_u(i), r_v(i)} \lim_{K \to \infty} \frac{1}{K} \sum_{i=0}^{K-1} [U(r_u(i)) + U(r_v(i))]$$
(24)

¹²The solutions were obtained for the following parameters: $C = 2.7 \cdot \mathbb{E}(D(i)), U(s(i)) = \log(1 + s(i)), \text{ and } \beta_{\text{nonlin}} = 1.3.$

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 $\{i, B_u(i), B_v(i)\}$,

$$h(i, B_u(i), B_v(i)) = \max \sum_{D_u, D_v} [U(r_u(i)) + U(r_v(i)) + h(i+1,$$

$$min[B_u(i) + q(D_u(i), B_u(i)) - s_u(i), C_u],$$

$$min[B_v(i) + q(D_v(i), B_v(i)) - s_v(i), C_v])],$$
(25)

where the maximization is over all $\{r_u(i), r_v(i)\}$ such that $c_{tx}r_u(i) + c_{rx}r_v(i) = s_u(i) \leq B_u(i)$, $c_{tx}r_v(i) + c_{rx}r_u(i) = s_v(i) \leq B_v(i)$. 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. 27 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 setup L-1 (see Fig. 10) is used as the random variable D_u and the random variable D_v .¹³

5.3.3 Approximate Policies and Heuristics

In this section, we examine low-complexity policies, and compare them to the discretized optimal policies (OPT). We consider the discretized model with $\mathcal{B}_h = \{0, \Delta, 2\Delta, ..., C\}$ and with D that takes values in the set $\{0, \Delta, 2\Delta, ..., D_{\max}\}$, where $D_{\max} \leq C$.

We consider Spend-What-You-Get (SG) policies, where the node aims to spend all the energy harvested in a slot, that is $s(i) \leftarrow Q(i) \forall i$. Similar policies were proposed in [55]. The SG policies have very low complexity. We additionally consider: Energy Storage Linear (SL) policies, where the spending rate is a linear function of the energy storage level B(i), that is, for some $0 \le \alpha \le 1$, $s_{\rm sl}(i) \leftarrow \alpha \cdot B(i)$, and Energy Storage Threshold (THR) policies, where the energy spending rate in a time slot i, $s_{\rm thr}(i)$, is assigned according to the energy storage level threshold under which the B(i)falls. Namely, $s_{\rm thr}(i) \leftarrow 0 \forall B(i) \le L_1$; $s_{\rm thr}(i) \leftarrow s_1 \forall L_1 < B(i) \le L_2$; ...; $s_{\rm thr}(i) \leftarrow s_T \forall B(i) > L_T$. Similar policies were proposed in [47].

We let π^*_{Δ} denote an optimal policy, π^{sg}_{Δ} denote an SG policy, and recall that \mathbb{E}_{π} denotes the

¹³The solutions were obtained for the following parameters: $C = 2.1 \cdot \mathbb{E}(D(i))$, $c_{tx} = \ln J/\text{bit}$, $c_{rx} = 2nJ/\text{bit}$, $U(s(i)) = \log(1 + s(i))$, and Q(i) = D(i) (linear energy storage).

expectation operator induced by policy π and the MDP model.

We first provide insights into the behavior of the discretized OPT policy (Observation 5) and the SG policy (Observation 6).

Observation 5 Under the OPT policy, it holds that $\lambda_{\Delta}^* \triangleq \lim_{K \to \infty} \mathbb{E}_{\pi^*}(Z/K) \leq U(\mathbb{E}(D)).$

Observation 6 Under the SG policy, it holds that $\lambda_{\Delta}^{sg} \triangleq \lim_{K \to \infty} \mathbb{E}_{\pi^{sg}}(Z/K) = \mathbb{E}(U(Q)).$

The performance of the SL and THR policies depends on the policy parameters. In the general case, the best policy parameters can be selected via "brute-force" algorithms. In special cases, the parameters may be selected using lower-complexity techniques, as we demonstrate next. Consider the set of SL policies which is identified with the set of possible values for the SL policy parameter. For a given parameter $0 \le \alpha \le 1$, the corresponding SL policy specifies the spending rate at slot *i* by $s(i) = \pi_{\rm sl}^{\alpha}(B(i)) = \lfloor \alpha \cdot B(i) \rfloor$. We next refer to an SL policy that maximizes the utility over all SL policies as an *optimal SL policy*. In a general scenario, the optimal parameter α can be computed by determining, for all feasible values of α , the state transition probabilities and the corresponding stationary state probabilities under the SL policy with the chosen α , and choosing the α that maximizes $Z^{\rm sl}$. Focusing on the uniform energy distribution, namely $p_d = \mathbb{P}\{D = d\} = 1/(D_{\rm max}+1)$ for all $d = 0, \Delta, ..., D_{\rm max}$, and on the linear energy storage model, we have the following:

Theorem 2 The optimal SL policy is given by $\pi_{sl}^*(b) = \lfloor D_{\max} \cdot (b/C) \rfloor$.

For example, suppose that C = 6, $D_{\text{max}} = 3$, and $\Delta = 1$. Then, the optimal SL policy is obtained with $\alpha = 3/6 = 0.5$. We note that in this case the SL policy coincides with the discretized optimal policy (computed numerically). The proof of Theorem 2 is provided in [22].

For link scenarios, similarly to the deterministic profile energy model, we can also use the DRC policies. In this case, the DRC policies are calculated using the marginal pdfs of D_u and D_v (rather than the joint pdf), and thus do not account for the possible dependency between D_u and D_v .

5.3.4 Numerical Results

We evaluate the performance of the OPT, THR, and SL policies via simulations.



Figure 28: Node scenarios with stochastic energy model: (a) objective function value, and (b) % node downtime.

The SL policies perform similarly to the OPT policies, and substantially outperform the THR policies, particularly when C is small. For example, Fig. 28 shows the performance of the THR1 (THR with one threshold), SL, and OPT policies. The policies were evaluated using, as an energy input, an empirical pdf of the diurnal energy recorded in L-1. Fig. 28(a) shows the upper bound on Z derived in Observation 5; the bound is tight when C is large. In these evaluations, the Z values under the SL and OPT policies are nearly identical. While under the SL and the OPT policies the downtimes are negligible, under the THR1 policy the nodes experience 1.8% -11.4% downtimes.

We additionally observe that in certain cases, the SL and the OPT policies coincide. For example, when C = 6, D = [0, 1, 2, 3], and $p_D = [0.25, 0.25, 0.25, 0.25]$, the OPT and the SL policies are identical.

Our study of the bounds on performance degradation due to quantization is motivated by the presence of such quantization in digital circuitry. In the EnHANTs prototypes, $\Delta = 8.6$ mJ (due to finite precision in the Coulomb counter that keeps track of the battery level, as we will explain in more detail in Section 6.1.1 in Chapter 6). It can be easily verified that for such Δ , in order for the results of Theorem 1 to apply, the maximum value of D should be at least 0.52 J (assuming uniform distribution of D). The typical D in our settings would require $\Delta \leq 0.3$ mJ. Obtaining bounds that apply to quantization parameters typically encountered in current energy harvesting devices (i.e., Δ that is same as, or larger, than the EnHANTs prototype Δ) is subject for future work.

5.4 Evaluating Model-free Approaches with Kinetic Energy Traces

In this section, we examine numerically the performance of a set of simple policies with kinetic energy traces we presented in Chapter 4. As we will demonstrate, kinetic energy traces are not well modeled by simple stochastic processes. Hence, rather than examining policies for stochastic energy models, we consider online policies (i.e., *model-free* approaches).

The long-term kinetic energy traces we use in this section are summarized in Table 5. In this section, we refer to trace identifiers M1, M2 listed in Table 5. Recall that we denote the power harvested by an inertial harvester by P(t). Using time slot length $T_{\text{int}} = 1$ second, we obtain a time-slotted process $P_{\text{meas}}(i)$ by calculating the average value of the P(t) for each T_{int} . We then calculate energy harvested, Q(i), as $Q(i) \leftarrow \eta \cdot P_{\text{meas}}(i) \cdot T_{\text{int}}$, where η is the conversion efficiency (assumed to be 20% [116]).

For the day-scale kinetic energy traces we collected, P_{meas} is *clearly not i.i.d. or Markov*. For example, for the P_{meas} for participant M1 and for $THR = 10 \ \mu\text{W}$, $p(P_{\text{meas}}(i) > THR | P_{\text{meas}}(i-1) > THR$, $P_{\text{meas}}(i-2) > THR$) = 0.91, while $p(P_{\text{meas}}(i) > THR | P_{\text{meas}}(i-1) > THR$, $P_{\text{meas}}(i-2) < THR$) = 0.54.

We use the following low-complexity online policies.

- Spend-what-you-get (SG): $s(i) \leftarrow Q(i)$.
- Storage-linear (SL): s(i) ← 2 · Q˜(i) · B(i)/C, where Q˜(i) is the running average of Q(i), that is, Q˜(i) ← ∑_{j=0}ⁱ⁻¹ Q(j)/i.
- Scheme-LB: s(i) ← (1 − ε) · Q̃(i) if B(i) + Q(i) ≥ (1 − ε) · Q̃(i), and s(i) ← B(i) + Q(i) otherwise, where ε is a small constant (we use ε = 0.01). The Scheme-LB policies were proposed and examined (for infinite C) in [11].

Examples of realizations of the SG, SL, and Scheme-LB policies with a kinetic energy trace are shown in Fig. 29. These realizations use the P_{meas} values for participant M5, whose P(t) was previously shown in Fig. 19(b). For these realizations, we use $C = 0.5 \cdot \sum_{i} P_{\text{meas}}(i)$. Fig. 29(a) shows the instantaneous data rates, r,¹⁴ while Fig. 29(b) shows the battery levels, B. The average

 $^{^{14}}$ For ease of visualization, Fig. 29(a) only shows r values in the 0–40 Kb/s range. This represents the full range of



Figure 29: Examples of realizations of the SG, SL, and Scheme-LB policies for participant M5: (a) data rates, r, and (b) battery levels, B.

data rate, \overline{r} , is 3.27 Kb/s under the SG policy, 2.95 Kb/s under the SL policy, and 3.07 Kb/s under the Scheme-LB policy.

We compare Z values for each policy to the maximal possible objective function value, $Z_{\text{max}} = \log[1 + (\sum_i e(i)/c_{\text{tx}})/K]$. As a measure of policies' energy neutrality, we consider the percentage of energy misused, E_{md} , which we define as $E_{\text{md}} = 100(\sum_i s(i)/\sum_i Q(i) - 1)$. A positive E_{md} represents overspending of the available energy and using some of the energy originally stored in the battery, and a negative E_{md} represents underspending of the energy. We also consider the node ON times. For these evaluations, for a time slot *i*, a node is ON if r(i) > 500 b/s and OFF otherwise.¹⁵

Policy Performance with Kinetic Energy Traces

The SL and the Scheme-LB policies perform well with the kinetic energy traces, particularly for large C values. For example, Fig. 30 shows the performance under the Scheme-LB, SL, and SG policies for a trace of participant M2. Fig. 30(a-d) show \overline{r} , Z (as a fraction of Z_{max}), ON time percentages, and E_{md} , respectively. These performance metrics are evaluated for different values of C, which are shown as a percentage of the total daily energy harvested, E_d . Each data point in Fig. 30(a-d) corresponds to a full policy realization, such as those shown in Fig. 29.

For relatively large C values (e.g. $C > 0.7E_d$), the SL and Scheme-LB policies perform well $-\overline{r}$

r values under the SL and Scheme-LB policies. Under the SG policies, instantaneous r values reach up to 180 Kb/s.

 $^{^{15}}$ For the evaluations with the kinetic energy traces, we use these performance metrics, rather than the node downtimes considered in the previous sections, because the processed kinetic energy values are often low, but are rarely exactly 0.



Figure 30: Performance of the Scheme-LB, SL, and SG policies for participant M2: (a) data rates, r, (b) objective function values, Z, (c) node ON time percentages, and (d) energy misuse percentages, $E_{\rm md}$.

is high (Fig. 30(a)), Z is close to Z_{max} (Fig. 30(b)), and ON times reach 100% (Fig. 30(c)). Both the SL and Scheme-LB policies slightly overspend the harvested energy (Fig. 30(d)), but the overuse does not exceed 5% of E_d .

Both the SL and the Scheme-LB policies dramatically outperform the SG policy. The SG policy achieves high \overline{r} (Fig. 30(a)) and fully uses the harvested energy (Fig. 30(d)), but performs poorly in terms of both objective function values and node ON times. Specifically, under the SG policy, Z/Z_{max} values are up to 2 times smaller than under the SL and the Scheme-LB policies (Fig. 30(b)), and the ON times are up to 4 times smaller (Fig. 30(c)). The poor performance under the SG policy emphasizes the need for energy management policies that consider motion energy variability and do not base spending rates directly on harvesting rates.

While the SL and Scheme-LB policies perform well, the performance differs from trace to trace. For example, Fig. 31 shows the performance of the SL policy with traces for participants M1, M2, and M5. Fig. 31(a,b) show the ON times and $E_{\rm md}$ values, correspondingly. It can be seen that in terms of both metrics, the SL policy performs better for M2 than for M1 and M5.



Figure 31: SL policy performance with traces for participants M1, M2, and M5: (a) ON times, and (b) misused energy percentages, $E_{\rm md}$.



Figure 32: Scheme-LB policy performance using energy traces (P_{meas}) for participant M1 and using the corresponding ON/OFF (P_{onoff}) , Markov (P_{markov}) , and i.i.d. (P_{iid}) processes: (a) average data rates, \overline{r} , and (b) node ON times.

Policy Performance with I.i.d. and Markov Processes

To assess the difference in performance between policies evaluated using real traces and i.i.d. and Markov processes, we use a different representation of the same energy harvesting process. Specifically, for a process P_{meas} calculated from our measurements, we generate an i.i.d. process, P_{iid} , by randomly permuting the values of P_{meas} . Recall that we defined an ON/OFF process P_{onoff} to be $P_{\text{onoff}}(i) \leftarrow 1$ ("ON") if P(i) > THR, and $P_{\text{onoff}}(i) \leftarrow 0$ ("OFF") otherwise (see Section 4.4.3 for more details). To generate a Markov process, P_{markov} , we first calculate the empirical state transition probabilities of the P_{onoff} process, $p_{0,1} = p(P_{\text{onoff}}(i) = 0|P_{\text{onoff}}(i-1) = 1)$ and $p_{1,0} = p(P_{\text{onoff}}(i) = 1|P_{\text{onoff}}(i-1) = 0)$. Then, we generate a Markov process with states $\{0, 1\}$ and transition probabilities $p_{0,1}, p_{1,0}$. We set the P_{markov} values for states 0 and 1 to the average values of $P_{\text{meas}}(i)$ for which $P_{\text{onoff}}(i) = 0$, and for which $P_{\text{onoff}}(i) = 1$, respectively. This ensures that the processes have the same first-order statistics.

The policy performance evaluated using i.i.d. and Markov processes differs dramatically from the

policy performance using the traces. For example, Fig. 32 shows \overline{r} and ON times obtained under the Scheme-LB policy for different processes based on a trace of participant M2. Using the process P_{onoff} , the performance is similar to the performance obtained using P_{meas} – the \overline{r} values differ by at most 17% (0.23 Kb/s), and the ON times differ by at most 7%. However, the performance evaluated using P_{markov} and P_{iid} differs greatly from the performance using P_{meas} . The differences in \overline{r} values reach over 105% (1.35 Kb/s), and the differences in ON times reach 63%.

Moreover, using Markov and i.i.d. processes results in different performance trends. Using $P_{\text{meas}}(i)$, the performance strongly depends on C, with \overline{r} for the different values of C differing by over 2.3 times, and with the ON percentages differing by over 45%. However, using P_{iid} and P_{markov} , both \overline{r} and ON times are nearly independent of C. Additionally, evaluating policy performance using P_{meas} demonstrates than ON times are an important metric – they can be relatively low for small values of C (Fig. 32(b)). However, when evaluating using P_{iid} and P_{markov} , the ON times are nearly 100% for all values of C, including values as low as 15% of the E_d . These results emphasize the need to evaluate energy harvesting-adaptive policies for wireless nodes equipped with an inertial harvester using real traces.

5.5 Conclusions and Future Work

In this chapter, we formulate resource allocation problems for energy harvesting nodes, aiming to allocate nodes' resources *in a uniform way with respect to time*. For the deterministic profile energy model and for the stochastic energy model, we formulate optimization problems and presents algorithms that determine energy and data rate allocation policies for single node and link scenarios. We also examine a set of simple policies, for some of which we obtain performance guarantees.

For the deterministic profile energy model, we demonstrate a set of algorithms of different complexities for different cases. That is, while for the most general settings the algorithms are relatively complex, in many settings we examined, optimal policies can be calculated with relatively low-complexity algorithms. For example, for node scenarios, we demonstrate relatively complex algorithms for general (non-linear) energy storage, simpler algorithms for the cases where energy storage is linear, and algorithms of very low complexity for the cases where energy storage is both large and linear. We used the algorithms to obtain numerical results for various cases, using, as energy inputs, light energy traces we described in Chapter 3. Additionally, for both deterministic profile and stochastic energy models, in many cases simple policies perform similarly to the optimal. For example, for deterministic profile energy model in link scenarios, the DRC algorithms perform well in many cases. For the stationary stochastic energy model, in the node scenarios the SL policies perform similarly to the optimal policies, and in certain cases coincide with the optimal.

Additionally, using kinetic energy traces we described in Chapter 4, we evaluate a set of simple online policies (i.e., *model-free* approach). Our evaluations demonstrate that in many cases simple policies perform well and emphasize the need to evaluate policies with real energy traces.

Future work may focus on addressing additional "working points" in the energy harvesting adaptive algorithm design space. In particular, energy harvesting adaptive algorithms for *partially predictable* environments and for energy harvesting *networks* may be subjects of future investigations. In Chapter 6, we evaluate some of the policies we have developed with energy harvesting hardware and light energy inputs based on the light energy traces.

5.A Proofs: Deterministic Energy Profile Environmental Energy Model

5.A.1 Single Node

Throughout this appendix, we define \hat{Q} to be $\hat{Q} \triangleq \sum_{i=0}^{K-1} Q(i)$, i.e., the total amount of energy harvested. We let $Q_{\max} = B_0 - B_K + \hat{Q}$ denote the maximum amount of energy a node can allocate. **Proof of Lemma 1**

Recall that functions U(s) used in the TFU problem are concave and non-decreasing (see Section 5.1). First notice the following facts.

Fact 1 The constraint sets of the TFU-LIN problem and the TFLA-LIN problem 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 the constraint set and the utility functions implies that any ϵ -decrease to one of the components of a solution to the TFU-LIN problem or a solution to the TFLA-LIN problem yields at most total ϵ -increase to the other components.

Fact 3 Let $y \le x$ and let f be a twice continuously differentiable concave function. Then for every $\epsilon > 0$, $f(x + \epsilon) + f(y - \epsilon) \le f(x) + f(y)$.¹⁶

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 \hat{s} , $\hat{s} \neq \bar{s}$, 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 components). By Fact 2, there exist j, k and $\epsilon > 0$, such that $\hat{s}(j) = \bar{s}(j) + \epsilon$ and $\hat{s}(k) = \bar{s}(k) - \epsilon$. 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, $U(\bar{s}(j)) + U(\bar{s}(k)) < U(\hat{s}(j)) + U(\hat{s}(k))$ where the second inequality is obtained by applying Fact 3. This contradicts the assumption that $\hat{s} \neq \bar{s}$, 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 \hat{s} is the (unique) solution to the TFU-LIN problem, it also the solution to the TFLA-LIN problem. Indeed, any ϵ -increase in one of \hat{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.

Proof of Lemma 2

Let s^* denote the energy allocation policy $s(i) = s^* = \hat{Q}/K \forall i$. First we show that the s^* -policy is feasible under the *LS Conditions* and the *LS-gen Conditions*. Then we show that, when feasible, the policy is the optimal solution to the TLFA-LIN problem.

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

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

Proof: Assume the energy storage capacity is sufficiently large. The storage state at the beginning of the *i*th slot under the *s*^{*}-policy 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, $B_0 \ge |\min_{1 \le i \le K} \tilde{B}(i)|$ is needed. In addition, in each time slot $i, C - B(i) \ge 0$ is required. Plugging in the expression for B(i), the condition $C - B_0 \ge \max_{1 \le i \le K} \tilde{B}(i)$ is obtained.

¹⁶Since f is concave and twice differentiable, f' is a decreasing function, and f'' < 0. Hence, using Taylor expansion, $f(x + \epsilon) + f(y - \epsilon) = f(x) + f'(x) \epsilon + f(y) + f'(y) (-\epsilon) + o(\epsilon^2) \le f(x) + f(y) + \epsilon(f'(x) - f'(y)) \le f(x) + f(y)$.

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

Proof: For any $\{Q(i)\}$ with a given $\sum_{i} Q(i) = H^*$, when the LS-gen Conditions hold, LS Conditions hold. This is easy to see by considering the vectors that obtain, over all $\{Q(i)\}$ with $\sum_{i} Q(i) = H^*$, $\min_{\{Q(i)\}} \tilde{B}$ and $\max_{\{Q(i)\}} \tilde{B}$. The $\min_{\{Q(i)\}} \tilde{B}$ is obtained for the $\{Q(i)\}$ such that $Q(i) = 0 \forall i < K - 1, Q(K - 1) = H^*$:

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

The $\max_{\{Q(i)\}} \tilde{B}$ is obtained for the $\{Q(i)\}$ such that $Q(0) = H^*$, $Q(i) = 0 \forall i > 0$:

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

Combining these two bounds with the LS Conditions, we obtain the LS-gen Conditions. Thus, when the LS-gen Conditions hold, the LS Conditions hold, and the s^* -policy is feasible by Proposition 1.

We now demonstrate that, when feasible, the s^* -policy is the optimal solution to the TFLA-LIN problem. Together, (5) and (7) imply that any other vector $\{s(i)\}$ in the feasible domain has at least one *i* such that $s(i) \leq s^*$. Thus the s^* -policy is lexicographically maximal, and therefore it is the optimal solution to the TFLA-LIN problem.

Proof of Lemma 3

Let s^* denote the energy allocation policy $s(i) = s^* = \hat{Q}/K \forall i$. The feasibility of the s^* -policy under the *LS Conditions* and the *LS-gen Conditions* is demonstrated in Appendix 5.A.1. When the s^* -policy is feasible, the TFU problem reduces to the following simple problem:

$$\max_{\{s(i)\}} \bar{U} = \sum_{i=0}^{K-1} U(s(i))$$
(26)

s.t.:
$$\sum_{i=0}^{K-1} s(i) \le \hat{Q}$$
 (27)

$$s(i) \ge 0 \ \forall \ i. \tag{28}$$

Below, we demonstrate that the s^* -policy solves this problem optimally for functions U(s(i)) satisfying conditions (I), (II), and (III) in Lemma 3.

U(s(i)) satisfying (I): since $\nabla \overline{U} \neq 0$ on the interior of the domain, the maximum is attained on the boundary. Since \overline{U} is symmetric with respect to s(i), it suffices to look at the boundaries of the form $\Omega_n = \left\{ s \in \mathbb{R}^K : \sum_{i=0}^n s(i) = \hat{Q}, s(i) \ge 0, s(j) = 0 \text{ for } j > n \right\}$ for $n = 1, \ldots, K - 1$. Applying the Lagrange multipliers method for Ω_n , we obtain $U'(s(i)) = \lambda$, $i = 0, \ldots, n$, which, due to the strict concavity of U(s(i)), implies local maximum of \overline{U} of the form $s(i) = \hat{Q}/n, i \le n, s(i) = 0, i > n$. From the concavity of U(s(i)), we may further conclude that the global maximum of \overline{U} is attained on Ω_{K-1} at the point $s(i) = \hat{Q}/K = s^* \forall i$.

U(s(i)) satisfying (II): The above analysis for U(s(i)) satisfying (I) can be applied here. This time we get only one point, $s(i) = s^* \forall i$, which is the global maximum.

U(s(i)) satisfying (III): Since the utility is being maximized, and since $U(s(i)) \to -\infty$ as $s(i) \to 0$, it is sufficient to solve (26) over a smaller region, that is, subject to (27) and to $s(i) \ge \epsilon$ for some $\epsilon > 0$ small. Following the method outlined above, the maximum is obtained at $s(i) = s^* \forall i$.

Proof of Observation 1

When the energy storage is sufficiently large, the optimal allocation is $s(i) = Q_{\max}/K \forall i$ (as demonstrated in [31]). The corresponding Z_K is $\tilde{Z}_K = \sum_i U(s(i)) = K \cdot U(Q_{\max}/K)$. When the energy storage is smaller, the total energy available to a node is at most Q_{\max} , and it is allocated less uniformly. Thus, due to the concavity of U, the Z_K for smaller storage conditions will be smaller, and therefore the above-stated \tilde{Z}_K is an upper bound. The result then follows by the monotonicity of U since $\sum_i Q(i) \leq \sum_i D(i)$.

Proof of Proposition 1

Denote \hat{Q}/K by A, and let $f = B_0/\hat{Q}$ (hence 0 < f < 1, and $B_0/K = f \cdot A$). From Observation 1, $Z_K^{\text{opt}} \leq K \cdot U(A)$. Since $\hat{Q} > B_0 = B_K$, there exists $s_{\text{cr}} \geq B_0/K$ such that if $s(i) \equiv s_{\text{cr}}$, a node spends (at least) B_0 units of energy and has sufficient B_K in storage at the end of the K slots. Hence, $Z_K^{\text{cr}} \geq K \cdot U(B_0/K) = K \cdot U(f \cdot A)$, and $Z_K^{\text{opt}}/Z_K^{\text{cr}} \leq (K \cdot U(A))/(K \cdot U(f \cdot A)) = (\log(1+A))/(\log(1+f \cdot A))$. For f = 1, $\log(1+A)/\log(1+f \cdot A) = 1$. For 0 < f < 1, this expression is a decreasing function of A, with its maximum achieved for $A \rightarrow 0$. Using l'Hôpital's rule, the following holds:

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

Hence, $Z_K^{\text{opt}}/Z_K^{\text{cr}} \leq 1/f = \hat{Q}/B_0$, and thus $Z_K^{\text{cr}} \geq Z_K^{\text{opt}} \cdot [B_0/\hat{Q}]$.

Proof of Proposition 2

For $s(i) = Q(i), Z_K^{sg} = \sum_i \log(M + Q(i)) = \sum_i \log(Q'(i)). \quad G(Q')$, the geometric mean of a sequence, can be transformed as $G(Q') \triangleq \sqrt[\kappa]{Q'(0) \cdot Q'(1) \cdot \ldots \cdot Q'(K-1)} = \exp((1/K) \cdot \sum_i \log(Q'(i))) = \exp((1/K) \cdot Z_K^{sg}).$ Thus $Z_K^{sg} = K \cdot \log(G(Q')).$ From Observation 1 we know that $Z_K^{opt} \leq KU(\sum_i Q(i)/K) = K \cdot \log(\overline{Q'}).$ Therefore, $Z_K^{sg}/Z_K^{opt} \geq [K \cdot \log(G(Q'))]/[K \cdot \log(\overline{Q'})],$ and consequently $Z_K^{sg} \geq Z_K^{opt} \cdot \log(G(Q'))/\log(\overline{Q'}).$

5.A.2 Link

Proof of Observation 2

Consider the case where the energy storage is unlimited. Then (4)-(8) reduce to $\sum_i s_u(i) \leq [B_{0,u} - B_{K,u} + \sum_i Q_u(i)], \sum_i s_v(i) \leq [B_{0,v} - B_{K,v} + \sum_i Q_v(i)]$. Since the objective function is concave and non-decreasing, and since the problem is symmetric for all slots *i*, the optimal solution will be obtained when each node assigns the same data rate to each slot. The same data rate for each slot will be obtained when each node assigns the same amount of energy to each slot, that is, $s_u(i) = [B_{0,u} - B_{K,u} + \sum_i Q_u(i)]/K \forall i, s_v(i) = [B_{0,v} - B_{K,v} + \sum_i Q_v(i)]/K \forall i$, and hence the overall solution will be $K \cdot \tilde{Z}_K$. This value is an upper-bound since for constrained energy storage conditions, the same amounts of energy will be assigned under additional constraints.

Proof of Observation 3

Notice that, similarly to Lemma 1, the constraint sets are the same for both the LTFL problem and the LTFU problem. When $c_{tx} = c_{rx} = \tilde{c}$, constraints (13), (14) are reduced to $\tilde{c} \cdot (r_u(i) + r_v(i)) \leq s_u(i)$ and $\tilde{c} \cdot (r_u(i) + r_v(i)) \leq s_v(i)$; combining them, we obtain $r_u(i) + r_v(i) \leq \min(s_u(i), s_v(i))/\tilde{c}$. Given the concavity of U(s(i)) in the LTFU problem, the utility is maximized when $r_u(i) = r_v(i) = \tilde{r}(i) = \frac{1}{2}\min(s_u(i), s_v(i))/\tilde{c}$. The optimal solution to the LTFL problem is max-min fair, thus it also assigns data rates as $r_u(i) = r_v(i) = \tilde{r}(i)$. Thus, the objective function of the LTFU problem can be stated as max $\sum_{i} U(\tilde{r}(i))$, and the objective function of the LTFL problem can be stated as

Lexicographically maximize $[\tilde{r}(0), ..., \tilde{r}(K-1)]$.

The equality of the optimal solutions to these problems under the same constraint sets follows from the proof of Lemma 1. \blacksquare

Proof of Lemma 4

When the LS or the LS-gen Conditions hold, constraints (4) – (8) for nodes u and v reduce to simple constraints on the sum of the energy spending rates: $\sum_{i=0}^{K-1} s_u(i) \leq \sum \hat{Q}_u$, $\sum_{i=0}^{K-1} s_v(i) \leq \sum \hat{Q}_v$. Therefore, solving, independently from each other, the TFLA problem, nodes u and vcalculate their energy spending allocations $\{s_u(i)\}, \{s_v(i)\}$ as $s_u(i) = s_u^* \forall i, s_v(i) = s_v^* \forall i$, where s_u^* and s_v^* are determined according to the node s^* -policy definition given in Lemma 2. Due to (17), the link data rates are assigned as $r_u(i) = r_v(i) = r^* = \min \{s_v^*, s_u^*\} / (c_{tx} + c_{rx})$ for each time slot i. Notice that any other vector $\{r_v(i), r_u(i)\}$ in the feasible domain has at least one i_0 such that $r_u(i_0) < r^*$ or $r_v(i_0) < r^*$. Thus the vector obtained by the LTFL-DRC algorithm is lexicographically maximal, i.e., the optimal solution to the LTFL problem.

Proof of Lemma 5

Similarly to the single node case (see Appendix 5.A.1), when LS or LS-gen Conditions hold, constraint sets (4) – (8) for nodes u and v reduce to simple constraints on the sum of the energy spending rates: $\sum_{i=0}^{K-1} s_u(i) \leq \sum \hat{Q}_u$, $\sum_{i=0}^{K-1} s_v(i) \leq \sum \hat{Q}_v$. Thus, by solving the TFU-LIN problem, the node energy spending rate vectors $\{s_u(i)\}$ and $\{s_v(i)\}$ are assigned as $s_u(i) = s_u^* \forall i$ and $s_v(i) = s_v^* \forall i$, where s_u^* and s_v^* are determined according to the node s^* -policy definition given in Lemma 3. Thus, in each slot *i*, the sub-problem we solve is the same, $\max_{r_u(i), r_v(i)} U(r_u(i)) + U(r_v(i))$ such that $c_{tx}r_u(i) + c_{tx}r_v(i) \leq s_u^*$, $c_{tx}r_v(i) + c_{tx}r_u(i) \leq s_v^*$, and the obtained data rate assignments are $r_u(i) = r_u^*$, $r_v(i) = r_v^* \forall i$. The optimality of these data rates for the LTFU problem follows directly from the arguments of the optimality of the s^* -policy in the proof of Lemma 3.

Proof of Proposition 3

The optimal solution to the LTFL is max – min fair, and thus, for each i, $r_u(i) = r_v(i) = r(i)$, and (13), can be reduced to $r(i) \leq (\min(s_u(i), s_v(i))/[c_{tx} + c_{rx}])$. Thus, the data rate in each slot *i* is fully determined by the minimum of $\{s_u(i), s_v(i)\}$. For $s_u(i) \leq s_v(i) \forall i$, the data rates are fully determined by the allocation of $s_u(i)$. The LTFL-DRC slot energy spending assignments $s_u(i)$ are lexicographically fair – a spending $s_u(i_1)$ cannot be increased without a decrease in some $s(i_2)$ that is already smaller. Thus, the energy allocation cannot be improved, and thus the LTFL-DRC solution is optimal.

Proof of Observation 4

Due to constraints (13), the data rates assigned by a DRC policy to a slot *i* will be zero if $s_u(i) = 0$ or $s_v(i) = 0$, thus $T_{L(u,v)}$ is not smaller than $\max[T_u, T_v]$. If both $s_u(i)$ and $s_v(i)$ are non-zero, then the r(i) values maximizing (18) will also be non-zero, thus $T_{L(u,v)}$ is not larger than $T_u + T_v$.

5.B Proofs: Stochastic Environmental Energy Model

Proof of Observation 5 The energy received in a slot *i* does not exceed D(i), and the overall expected amount does not exceed $K \cdot \mathbb{E}(D)$. The concave objective function *U* is maximized when the energy is spent uniformly, thus for the total expected energy $K \cdot \mathbb{E}(D)$, the utility is maximized for energy spending rate $s(i) = [K \cdot \mathbb{E}(D)]/K = \mathbb{E}(D) \forall i$. Hence, $\mathbb{E}(Z_{opt})$ is bounded as $K \cdot (1/K) \cdot U(\mathbb{E}(D)) = U(\mathbb{E}(D))$.

Proof of Observation 6 $\mathbb{E}(Z_{sg}) = \lim_{K \to \infty} \frac{1}{K} \sum_i U(Q(i)) = \mathbb{E}(U(Q)).$

Chapter 6

EnHANT Prototypes Testbed for Energy Harvesting Adaptive Policy Evaluations

In this chapter, we describe the design and development of Energy Harvesting Active Networked Tags (EnHANTs) prototypes and the EnHANTs prototypes testbed. The prototypes and the testbed were developed over the past 4 years in 6 integration phases. At the end of each phase, we presented the prototypes and the testbed in a conference demonstration session [26,28,57,80,85,119]. The current prototypes (Fig. 33(a)), are larger than the envisioned EnHANTs (show previously in Fig. 4). Yet, the prototypes already harvest indoor light energy using custom-designed organic solar cells which can be made flexible, and communicate wirelessly using ultra-low-power Ultra-wideband Impulse-Radio (UWB-IR) transceivers. The developed EnHANTs testbed enables controllable and repeatable experiments with communications and networking algorithms for energy harvesting nodes. The testbed allows observing the states of the prototypes in real time. It also includes a *software-based light control system* (Fig. 33(b)) that can expose the prototypes to controllable light conditions based on real-world light energy traces. For instance, it can "replay" the indoor light energy traces we presented in Chapter 3.



Figure 33: Current EnHANTs prototype testbed: (a) current EnHANT prototype, and (b) a software-based light control system, along with 4 EnHANT prototypes.

The prototypes form small networks and adapt their communications and networking patterns to energy harvesting states. We implemented energy harvesting adaptive optimal and approximate policies we examined analytically in Chapter 5. We evaluated some of these policies with the help of the developed testbed functionalities. We also implemented energy harvesting adaptive network algorithms (e.g., flow control, topology adaptations) in the prototypes. We demonstrate the performance of simple policies for flow control and collection tree adaptations. To the best of our knowledge, this work is the first attempt to evaluate energy harvesting adaptive policies in a controllable experimental environment.

The design and development of the EnHANTs are a joint effort of several research groups. Our contributions are in the design and development of energy harvesting adaptive algorithms, energy harvesting module interface design and integration, overall prototype integration, and the design and integration of the testbed control systems' interfaces. Major design contributions were also made by Ph.D. candidates J. Sarik and B. Vigraham (hardware) and Ph.D. candidate R. Margolies (medium access control, transceiver integration, prototype integration). Versions of different prototype and testbed components were designed as part of undergraduate and M.S. student projects (see Appendix A), as acknowledged via authorship in [26, 28, 57, 80, 85, 119]. Projects of G. Stanje, E. Katz, D. Roggensinger, and H. Huang contributed to the energy harvesting adaptive networking functionalities of the EnHANTs prototypes.

In this chapter, we first overview the EnHANTs prototypes and testbed functionalities and design (Section 6.1). We then present the results of evaluations of some of the policies we analyzed in Chapter 5 (Section 6.2). Finally, we present evaluation results for energy harvesting adaptive



Figure 34: A block diagram of the EnHANT prototype, and its interactions with the testbed.

EnHANTs networking policies (Section 6.3). A brief summary of the functionalities implemented in the different EnHANTs prototype integration phases is provided in Appendix 6.A.

We previously described the high-level EnHANTs design in [23, 32] and presented the EnHANT prototype and testbed design and development in [25].

6.1 EnHANT Prototypes and Prototype Testbed

In this section, we overview the EnHANT prototypes and the prototype testbed we developed. We focus on the prototype and testbed functionalities that are essential for energy harvesting adaptive policy evaluations. The other prototype and testbed features are described in more detail in [24, 25].

6.1.1 EnHANT Prototype

The EnHANT prototype is shown in Fig. 33(a). Each prototype includes an Energy Harvesting Module interfaced with a solar cell, a UWB-IR Communication Module, and a Control Module. The prototype block diagram, including the different modules and their interactions, is shown in Fig. 34.

- Energy Harvesting Module (EHM) The EHM, described in detail below, contains a rechargeable battery (used to store the energy harvested by a solar cell) and energy monitoring circuitry, and provides real time energy harvesting awareness.
- Communication Module The prototypes communicate with each other wirelessly using



Figure 35: Energy Harvesting Module (EHM): (a) a photo, and (b) a block diagram.



Figure 36: Solar cells integrated with EnHANT prototypes: (a) an amorphous silicon (a-Si) solar cell, and (b) a custom-fabricated organic photovoltaic (OPV).

ultra-low-power UWB-IR Communication Modules, based on UWB-IR transmitter and receiver chips previously described in [15]. The custom chips are mounted onto a printed circuit board that interfaces with the other prototype components. A Complex Programmable Logic Device is used to realize the "glue logic" between the radio chipset and the rest of the prototype.

• Control Module – Based on a legacy off-the-shelf MICA2 mote, the Control Module runs TinyOS with an added Fennec Fox software framework [87]. The Control Module is integrated with the Communication Module such that packets originating in the TinyOS application layer are sent wirelessly via the UWB-IR transceiver. The control module implements the medium access control and forwarding protocols tailored for the UWB-IR transceivers. More detailed description of the UWB-IR protocols we employ is available in [24, 25].

The prototypes adapt their networking and communication patterns based on the energy states, which are monitored by the EHM. A block diagram and a photo of the EHM are shown in Fig. 35. The EHM monitors the battery level, B(i), and the energy harvesting rate, e(i), and reports them to the Control Module. To track B(i), the EHM's energy monitoring circuitry uses a Coulomb counter, which measures the bidirectional current across R_{SENSE2} . To track e(i), the EHM uses a high side current sense amplifier, which measures the instantaneous current across R_{SENSE1} . The Coulomb counter updates the battery level every 0.875 seconds; the resolution in B(i) is under 5 mC. This allows for tracking the energy storage level in *nearly real time*.

The EHM does not supply energy to the other EnHANT prototype components. Rather, as shown schematically in Fig. 34, the EHM implements *controlled energy spending functionality*. Specifically, in correspondence with transceiver energy spending on transmitting and receiving packets, the Control Module signals to the EHM to activate a small load, which spends energy at a requested rate s(i). Releasing the constraint of running the prototype using harvested energy allows us to experiment with various hardware and protocol configurations. In the EHM, the energy is spent by discharging the battery through the load resistor R_{LOAD} (see Fig. 35) for τ ms. Each load activation reduces the battery by $\Delta B = \tau \cdot V_{\text{BAT}}/(R_{\text{LOAD}} + R_{\text{SENSE1}}) = 208.68 \ \mu\text{C}$. To verify the precision of the EHM's controllable energy spending, we compared the energy spending rates calculated according to this formula and the energy spending rates we experimentally obtained, for a set of EHM load activation rates. For up to 12 load activations per second, the discrepancy was under 2.1%.

For energy storage, the EHM uses a 2.4 V NiMH battery with 150 mAh (1,296 J) capacity. The EHM battery is intentionally oversized. This allows us to conduct experiments with different values of battery capacity C by restricting (in software) the battery operating range.

We equipped EnHANT prototypes with custom-designed organic photovoltaics (OPVs) and with commercially available amorphous silicon (a-Si) solar cells that are commonly used for indoor light energy harvesting applications [100,115]. The custom OPVs and the commercial a-Si solar cells are shown in Fig. 36. The a-Si cells are the Sanyo AM-1815 cells with a 5.61x4.52 cm² active area. A description of the process we followed to fabricate the OPVs can be found in [98]. The OPVs area is 5.0x5.0 cm². The two types of photovoltaics can be easily interchanged, as can be seen in Fig. 33(a).

6.1.2 EnHANT Prototype Testbed

The small-scale EnHANTs testbed includes a *control and monitoring system* and a *software-based light control system*. The testbed is shown schematically in Fig. 37. The current testbed can include up to 6 EnHANT prototypes.



Figure 37: A schematic diagram of the EnHANTs testbed.



Figure 38: A screenshot of the EnHANTs testbed monitoring system.

Prototype Control and Monitoring System – For control and monitoring, the prototypes are placed on MIB600 programming boards and accessed from a PC via Ethernet. On the PC, a Java-based graphical monitoring system records and shows in real time data rates, r, energy harvested, e, and battery levels, B, of each of the prototypes, and the individual packets transmitted (they are shown as flashing "arrows"). A screenshot of the monitoring system is shown in Fig. 38. **Software-based Light Control System** – The software-based light control system allows exposing individual prototypes to repeatable light energy conditions *based on real-world irradiance traces*. The system was previously shown in Fig. 33(b). To ensure full control over light conditions, the prototypes' solar cells are placed inside custom-designed dark box enclosures, as shown in Fig. 39. The light sources are component cool white Light-Emitting Diodes, mounted on heat sinks and attached to the enclosures. A LabVIEW script on a PC controls the irradiance inside each enclosure by controlling the current supplied to the Light-Emitting Diodes from a DC power supply. The irradiance produced by the system was calibrated using a NIST-traceable photodiode (Newport UV-818). The system can produce over 3,000 distinct irradiance levels between 0 and 14 mW/cm²



Figure 39: A dark box enclosure used in the EnHANTs testbed software-based light control system.



Figure 40: Energy harvesting for a prototype with an a-Si solar cell and with an organic photovoltaic (OPV): (a) generated power as a function of irradiance, (b) time-varying irradiance based on a light energy trace, and (c) power harvested by an a-Si solar cell and an OPV exposed to this irradiance.

(an effective resolution of less than 5 μ W/cm²); the irradiance levels can be changed with time steps of under 0.1 second.

With the software-based light control setup, we extensively use the *indoor irradiance traces* we presented in Chapter 3, *exposing the prototypes to the light conditions based on the traces*. We refer to the traces by their identifiers, L-1, L-2, ... (see Table 1 in Chapter 3).

The software-based light control system allows replicating real-world irradiance traces with remarkable repeatability (e.g., the energy harvesting rates over *multiple repetitions* of light energy traces are shown in Fig. 42(a), 43(a), 47(a), 48(a)). This ensures that experimental evaluations of energy harvesting adaptive policies are *based on the same energy inputs*. To be able to conduct many experiments with trace inputs in a reasonable amount of time, we "compress" (downsample) day-long traces. To capture the corresponding dynamics in the energy storage behavior, we also scale the light levels (by a factor indicated). While primarily designed for evaluating energy harvesting adaptive communications and networking policies, the software-based light control system can also be used for other evaluations. For example, using this system, we demonstrated that *different types of solar cells perform differently under the same light conditions*. Fig. 40 demonstrates an example of evaluations of different solar cells, an OPV and an a-Si cell. The harvesting efficiency of the OPVs is 1%, while the harvesting efficiency of a-Si cells varies between 1% and 3% depending on the irradiance (see the measurement results in Fig. 40(a)). Correspondingly, as shown in Fig. 40(c), when we expose the two solar cells to irradiance levels based on a light energy trace recorded over a day in location L-1 (Fig. 40(b)), *the "curves" of the power generated by the two solar cells have different shapes*. We note that these effects are difficult to capture in simulations, which simply assume that the energy harvested by a solar cell is a linear function of the irradiance (e.g., [19, 40]).

6.2 Evaluating Energy Harvesting Adaptive Policies

In this section, we evaluate experimentally the policies we examined analytically for the *deterministic profile* energy model in Chapter 5. We present evaluation results for energy harvesting *nodes* (Sections 6.2.1) and *links* (Section 6.2.2). In Section 6.2.3 we additionally briefly comment on the evaluation of policies we examined analytically for node and link scenarios for the *stochastic* energy model. We describe policies we have implemented and examined for *networks* of energy harvesting nodes in a subsequent Section 6.3.

We conducted extensive experiments with different policies, providing nodes, equipped with either an a-Si solar cell or an OPV, a dynamic light energy input based on the indoor light energy traces summarized in Chapter 3. Each EnHANT prototype monitors its energy harvesting rate, Q, battery level, B, and energy spending rate, s. The prototypes exchange these values via wireless communications, and adapt their data rates, r, according to the policies implemented. We report Bin either *Coulombs* or *Joules*.¹ We record the prototype parameters via EnHANTs testbed control and monitoring functionality. For the prototypes, the energy cost to transmit is $c_{tx} = 35.5$ nJ/bit,

¹In general, battery levels and battery capacities are often reported in various different units, units of charge (Coulombs, Ampere-hours), and units of energy (Joules). Since $1 \text{ C} = 1 \text{ A} \cdot 1$ second, the conversion between Coulombs and Ampere-hours is straightforward. To calculate the *energy*, the charge is multiplied by the voltage of the battery, V_{BAT} . In our system, $V_{\text{BAT}} = 2.4 \text{ V}$, thus the charge of 1 C (or, equivalently, 278 μ Ah) corresponds to 2.4 J of energy.



Figure 41: Sample EnHANT prototype policy realizations: (a) OPT, SG, and CR policies for a node scenario with a deterministic profile energy model, and (b) OPT, DRC-NOPT, DRC-SG, and DRC-CR policies for a link scenario with a deterministic profile energy model.

and the energy cost to receive is $c_{\rm rx} = 384.5 \text{ nJ/bit } [24, 25].$

This section uses the notation, model, and policies we introduced in Chapter 5. Similarly to Chapter 5, we denote the time slot index by *i*. The prototypes create their energy profiles $\{\overline{Q}(1), ..., \overline{Q}(K)\}$, used in most policies we examine for the deterministic profile energy model, by determining their expected harvesting rates, $\overline{Q}(i)$, for the different time intervals. In node scenarios, we evaluate the Optimal (OPT), Constant Rate (CR), and Spend-what-you-get (SG) policies, introduced in Section 5.2.1. Sample EnHANT prototype realizations of these policies are shown in Fig. 41(a).² In link scenarios, we evaluate the OPT policies and several variants of the Decoupled Rate Control (DRC) policies, introduced in Section 5.2.2. Specifically, we examine Node-optimal DRC (DRC-NOPT), Constant Rate DRC (DRC-CR), and Spend-what-you-get DRC (DRC-SG) policies. Fig. 41(b) demonstrates the (overlapping) $r_u(i)$ and $r_v(i)$ values for a set of example EnHANT prototype realizations of these policies.³

Recall that in Chapter 5 we considered *linear* and *general* energy storage types, corresponding to batteries and capacitors. EnHANT prototypes store the energy in a battery. The experiments presented in this section thus correspond to the *linear energy storage*. As battery-based and capacitorbased systems use different circuitry [120], non-trivial circuitry redesign effort would be required to adapt the prototypes to use capacitors.

²These realizations were obtained for C = 0.069 J and with the light energy input corresponding to a day in L-3, scaled by 22x.

³These realizations were obtained for $C_u = C_v = 0.069$ J, and with the light energy input corresponding to a day in L-1 and L-2, scaled by 104x.



Figure 42: Node scenarios with a deterministic profile energy model: (a) energy harvesting rates, (b) energy spending rates and (c) battery levels, experimental and obtained via simulations, and (d) objective function values under the OPT and CR policies.

6.2.1 Evaluating Optimal and Approximate Node Policies

As mentioned above, in node scenarios, we evaluate the OPT, CR, and SG policies.

Prior to the testbed implementations and evaluations, we have extensively evaluated our policies via MATLAB-based simulations. In nodes scenarios we examined, the results obtained experimentally and via MATLAB-based simulations correspond closely. For example, Fig. 42(b) and Fig. 42(c) show the s(i) and B(i) obtained under the OPT policies, for two different values of C, $C_1 = 0.14$ J and $C_2 = 0.09$ J. For these experiments, we provided the prototype with a light input based on the light energy trace recorded over a day in L-3 (scaled by 22x). The energy harvesting rates Q(i)are shown in Fig. 42(a); the errorbars demonstrate the variations in Q(i) in different experiments. The s(i) "curves" (Fig. 42(b)) differ, in the simulations and in the experiments, only by a short delay, caused by the inexact correspondence of time slots in the simulations and in the experimental results. The B(i) "curves" (Fig. 42(c)) correspond closely, with only small discrepancies arising, due to the difference in B(i) update intervals, when B(i) is close to the energy storage capacity C.
This precise correspondence confirms the reliability and precision of the energy state monitoring and controlled energy spending functionalities of the EHM.

The performance of the CR policies depends on the energy storage capacity C. When C is small, the OPT policies outperform the CR policies. When C is large, the performance of the OPT and CR policies is similar (as we described in Chapter 5, under certain conditions the OPT and CR policies coincide). This can be seen, for example, in Fig. 42(d) which shows the objective function values, Z, under different policies for different values of C, for the Q(i) shown in Fig. 42(a). Each data point in Fig. 42(d) corresponds to a separate experiment, repeated 3 times.

The CR policies outperform the SG policies. In the experiments described above, for example, the Z^{sg} for all values of C is only 86.0% of the Z^{cr} obtained with the smallest C we considered. Additionally, under the SG policies, the nodes experience substantial downtimes. For example, for L-1 – L-4, the node downtimes under the SG policies are 22 – 52%. Finally, the upper bound on Z, demonstrated in Observation 1 in Section 5.2.1, is tight when C is large, as can be seen, for example, in Fig. 42(d).

6.2.2 Evaluating Optimal and Approximate Link Policies

As mentioned above, we evaluate the OPT policies and several variants of the DRC policies: DRC-NOPT, DRC-CR, and DRC-SG. For the evaluations, we use light energy traces *concurrently recorded in nearby locations*. In the link scenarios, the experimentally obtained data rates are slightly lower than the simulated data rates (as can be seen, for example, in Fig. 43(c,d)). This is due to packet errors that are not reflected in our simulations, and are the result of packet collisions, system computational overloads, and noise-induced bit errors.

The performance of the DRC-NOPT policies is similar to the performance of the OPT policies, the performance of the DRC-CR policies is dependent on C, and the DRC-SG policies are outperformed by the other policies.

Example evaluation results are shown in Fig. 43. We provided the prototypes with light inputs corresponding to the light energy traces recorded over a day in L-1 (for node u) and L-2 (for node v), scaled by 110x. The nodes' Q(i) values in these experiments are shown in Fig. 43(a) (with the errorbars corresponding to the variations in Q(i) between the different repetitions). The Z



Figure 43: Link scenarios with a deterministic profile energy model: (a) energy harvesting rates, (b) objective function values, (c) average data rates under the DRC-CR policies, and (d) average data rates under the DRC-NOPT policies.

values under the different policies are shown in Fig. 43(b). The r values are shown in Fig. 43(c,d). Each data point in Fig. 43(b-d) corresponds to a separate experiment, repeated 5 times. In these experiments, for all values of C, the $Z^{drc-nopt}$ is at least 99.1% of the Z^{opt} . Additionally, the upper bound on Z, derived in Observation 2 in Section 5.2.2, is tight when C is large (see, for example, Fig. 43(b)).

The DRC-SG policies result in substantial downtimes (while under the other policies no link downtimes are experienced). For example, for a link with nodes' energy inputs corresponding to (L-1, L-2), $T_{L(u,v)} = 57\%$, and for (L-2, L-3), $T_{L(u,v)} = 64\%$. These downtimes are close to the lower bound on $T_{L(u,v)}$ derived in Observation 4 in Section 5.2.2.

6.2.3 Policy Implementations for the Stochastic Energy Model

The experimental results presented in this work are restricted to the *deterministic profile* environmental energy model. For the *stochastic* energy model, we implemented the OPT, Storage-statelinear (SL), and Threshold-based (THR) policies in the EnHANTs prototypes. Sample EnHANT



Figure 44: Sample EnHANT prototype policy realizations for a node scenario with a stochastic energy model: OPT, SL, and THR1 policies.



Figure 45: Small-scale multihop network topologies: (a) a 3-node *line* network, and (b) a 4-node *diamond* network.

prototype realizations of these policies are demonstrated in Fig. 44, for example (these realizations were obtained for C = 0.4313 J; we use light energy input that corresponds to the pdf of the diurnal energy in L-2).

Due to the relatively slow prototype energy storage state changes when charging and discharging, for these policies, the full experimental evaluations (which require a large number of samples) are subject for future work. The close correspondence between the simulations and the testbed experiments for the policies with deterministic energy inputs (e.g., Fig. 42(b,c)) suggests that these evaluations will closely correspond to the simulations as well.

6.3 Energy Harvesting Adaptive EnHANT Networking

In this section, we evaluate, using a small multihop network of EnHANT prototypes, policies for *networks* of energy harvesting nodes. We focus on data collection scenarios (e.g., ID collection), corresponding to the tracking applications we envision for the EnHANTs. Fig. 45 shows the considered network topologies: a 3-node *line* network and a 4-node *diamond* network. In these topologies,



Figure 46: Node energy allocation policies used as building blocks for network policies: (a) energy harvesting rates, (b) energy spending rates under the EX, EP-1, and EP-12 policies, and (c) battery levels, experimentally measured and simulated, under the EX and EP-1 policies.

the prototypes u, v, and w generate messages, and send them, via multihop collection trees, to a prototype that serves as a Collection Coordinator (CC). We note that most algorithms proposed for networks of energy harvesting nodes [10,40,55] are too complex for implementation in ultra-lowpower indoor environments, as they require multiple local [10] or global [55] iterations, or complex calculations [40]. We thus focus on simple policies.

As building blocks for wider-scale (i.e., network) energy harvesting adaptive policies, we use the following node energy adaptive policies that closely match node energy spending rates, s(i), to node energy harvesting rates, e(i).

- Exponential policies (EX) The desired energy spending rate s(i) is set to the exponential average of the energy harvesting rate: s(i) ← ê(i) = α · ê(i − 1) + (1 − α) · e(i), 0 ≤ α ≤ 1. Similar policies were evaluated, via simulations, in [55].
- Energy Profile-based policies (EP-K) In the EP-K policies (K corresponds to the number of time intervals), the node's desired energy spending rates are set to the expected energy harvesting rates: s(i) ← ē(i) ∀i ∈ K. For example, the EP-1 policy (examined, via simulations, in [19]), corresponds to a node spending energy at its average expected harvesting rate over the entire planning horizon (for large C, EP-1 and CR policies coincide).

Examples of energy spending rates, recorded in EnHANT prototypes running these policies, are shown in Fig. 46(b). They are obtained for node energy harvesting rates illustrated in Fig. 46(a) (where errorbars represent variations in energy harvesting rates in different experiments).⁴ The EX

 $^{^{4}}$ We provided the prototype with a light input corresponding to the light energy recorded over a day in L-3,



Figure 47: FLEX flow control policies: (a) energy harvesting rates, (b) data rates under the EX and EP-1 node energy allocation policies, and (c) average data rates, under the EX policy, for different values of r_c .

and EP-K policies effectively ensure *energy neutrality* (see Section 5.1), that is, effectively match energy spending to energy harvesting. For example, in all our experiments with the EX and EP-K policies corresponding to node energy harvesting rates shown in Fig. 46(a), the node used 95%-96%of the harvested energy.

6.3.1 Flow Control Policies

Our experiments with networks of energy harvesting nodes strongly indicate the need for flow control policies. When nodes set their data rates without considering other nodes (i.e., by setting $r(i) \leftarrow s(i)/c_{tx}$), networked nodes overspend energy dramatically. In the UWB-IR-based EnHANT prototypes, this is particularly pronounced since, due to the c_{tx}/c_{rx} ratio, a prototype spends approximately 10 times more energy on receiving and forwarding a packet for another prototype than on transmitting its own packet.

We implemented and evaluated flow control policies to which we refer to as *FLEX*, that are based on the *DLEX* node data rate assignment algorithm proposed in [19]. FLEX, running on the CC, assigns fair (*lexicographically maximal* [19]) data rates to the network nodes.⁵ Under FLEX, the prototypes independently determine their desired energy spending rates, s(i), using the EX or EP-K node energy allocation policies, and send them to the CC. The CC allocates data rates such that the total energy spending rates of the nodes do not exceed s(i), and the assigned data rates are

compressed to 321 seconds and scaled by 2.1x. The prototype was equipped with an a-Si solar cell. In the experiments conducted, the variability in total energy harvested was *under 1.9%*.

 $^{^{5}}$ The DLEX data rate allocation algorithm developed in [19] is implicitly tied to a particular, EP-based, node energy allocation policy. In FLEX, we combine the data rate allocation algorithm of [19] with different node policies.

Algorithm 3 FLEX policy running on the CC, for the 3-node *line* multihop network topology shown in Fig. 45(a).

 $\begin{array}{l} \textbf{Input:} \ s_u(i), s_v(i), s_{cc}(i); \\ \textbf{if} \ s_u(i) > s_v(i) \cdot (2 + c_{\text{rx}}/c_{\text{tx}}) \textbf{ then} \\ \widehat{s}_v(i) \leftarrow s_v(i); \ \widehat{s}_u(i) \leftarrow s_u(i) - s_v(i) \cdot (1 + c_{\text{rx}}/c_{\text{tx}}); \\ \textbf{else} \\ \{ \widehat{s}_u(i), \widehat{s}_v(i) \} \leftarrow s_u(i)/(2 + c_{\text{rx}}/c_{\text{tx}}); \\ \textbf{if} \ (s_u(i)/c_{\text{tx}}) > (s_{cc}(i)/c_{\text{rx}}) \textbf{ then} \\ R_T \leftarrow s_{cc}(i)/c_{\text{rx}}; R_{\text{curr}} \leftarrow [\widehat{s}_u(i) + \widehat{s}_v(i)]/c_{\text{tx}}; \\ \textbf{if} \ R_{\text{curr}} > R_T \textbf{ then} \\ \widehat{s}_u(i) \leftarrow \widehat{s}_u(i) \cdot R_T/R_{\text{curr}}; \ \widehat{s}_v(i) \leftarrow \widehat{s}_v(i) \cdot R_T/R_{\text{curr}}; \\ \textbf{Return:} \ r_u(i) \leftarrow \max[\widehat{s}_u(i)/c_{\text{tx}}, r_{\text{min}}]; \\ r_v(i) \leftarrow \max[\widehat{s}_v(i)/c_{\text{tx}}, r_{\text{min}}]; \end{array}$

fair. The FLEX policy for the 3-node line network topology (Fig. 45(a)) is shown in Algorithm 3. The algorithm first computes lexicographically maximal data rates that the forwarding node u can support, then checks whether the CC can support these rates, and, if necessary, scales the rates proportionally. To maintain network connectivity, nodes communicate at a rate of at least r_{\min} .

We conducted extensive experiments with the FLEX policy using a variety of light inputs. Fig. 47(b), for example, demonstrates data rates assigned by FLEX, in combination with the EX and EP-1 policies, for a network with node energy harvesting rates shown in Fig. 47(a) (where errorbars represent variations in energy harvesting rates in different experiments).⁶

The FLEX policy, in combination with the EP-1, ensures energy neutrality. In combination with EX and with EP-K for $K \neq 1$, FLEX may underspend or overspend the energy of the nodes. For example, in the evaluation scenarios shown in Fig. 47, combined with EP-1, FLEX spends 96.4% of the energy harvested by node u; combined with EX, it spends 125%. To achieve energy neutrality, FLEX needs to take into account energy spending on control messages, transmitted by the CC at a fixed rate r_c . For example, prior to calculating node data rates according to Algorithm 3, nodes' s(i) values need to be reduced as $s_{cc}(i) \leftarrow s_{cc}(i) - r_c \cdot c_{tx}$, $s_u(i) \leftarrow s_u(i) - r_c \cdot [c_{rx} + c_{tx}]$, $s_v(i) \leftarrow s_v(i) - r_c \cdot c_{rx}$. The r_c values directly affect nodes' data rates, as can be seen, for example, in Fig. 47(c), where each data point corresponds to a complete 11.6 minute (700 second) testbed

⁶We provided the prototypes with light inputs corresponding to the light energy recorded over a day in L-2, compressed to 720 seconds (12 minutes) and scaled by 60.4x. The prototypes were equipped with OPVs. In 8 experiments, the variability in the total energy harvested was under 4.1% for the CC, and under 1% for nodes u and v. We provided nodes with nearly identical light inputs, yet harvesting rates, shown in Fig. 47(a), differed by more than 1.8x. This is due to different efficiencies of the OPVs integrated with different prototypes.



Figure 48: Topology adaptation policies: (a) energy harvesting rates, and (b) percentage of harvested energy used by the nodes.

Table 7: Data rates \overline{r} under different collection tree adaptation policies.

Table 1. Data fates 7 ander unterent concerton tree adaptation ponetes.									
Policy	Tree A	Tree B	RR- $T, T = 10$	RR- $T, T = 50$	L- $B', B' = 7$	L- $B', B' = 18$			
\overline{r} , % of MX	82.4	95.1	92.6	97.8	88.2	92.8			

experiment with the energy harvesting rates shown in Fig. 47(a).

6.3.2 Collection Tree Adaptation Policies

We implemented and examined a number of collection tree adaptation policies in a *diamond* network topology, shown in Fig. 45(b). In this topology, nodes u and v send messages to the CC directly, while node w sends its messages via a forwarder, u (*tree A*) or v (*tree B*). A collection tree is chosen by the CC based on one of the heuristics outlined below. Once a tree is selected, node data rates are assigned using the FLEX flow control policies, in combination with the EX node energy allocation policies.

We compare the performance of the following policies to the performance in networks where the collection trees are fixed, and to an EP-1-based *MX policy*, that calculates the best collection tree offline.

- Round Robin (RR-T) Trees A and B alternate every T seconds. We consider the RR-T policies with T = 10 seconds and T = 50 seconds.
- Battery Level-based (L-B') The collection tree is changed if the battery level of the forwarder in the current tree is B' mC lower than the battery level of the other forwarder. We consider the L-B' policy with B' = 7 mC and B' = 18 mC. We note that tree selection

based on node battery levels is commonly used in traditional (non-energy-harvesting) sensor networks [109].

Policy evaluation results presented below correspond to the node energy harvesting rates shown in Fig. 48(a), where errorbars represent variations in harvesting rates in different experiments.⁷ Table 7 shows the average data rates obtained under different policies, as a percentage of data rates obtained under the MX policy. Note that using Trees A and B results in different data rates (due to the difference in the energy harvested by u and v). Fig. 48(b) shows the percentage of harvested energy used by the nodes.

All collection tree adaptation policies resulted in data rate improvements over the fixed Tree A. This emphasizes the need for *topology adaptation policies* for networks of energy harvesting nodes. The use of the L-B' policies *did not result in data rate improvements* over the fixed Tree B. Due to FLEX and EX not taking battery levels into account, L-B' only changes energy use patterns in nodes u and v (as can be seen in Fig. 48(b)).

6.4 Conclusions

In this chapter, we present the design considerations for EnHANTs prototypes and prototypes testbed. The prototypes communicate with each other via custom-developed UWB-IR transceivers and harvest indoor light energy using custom-designed organic solar cells. They form small multihop networks and adapt to the energy harvesting states in real time. The developed EnHANTs prototype testbed is the first system that allows exposing energy harvesting nodes to repeatable light energy conditions based on real-world irradiance (light energy) traces.

We used the testbed to evaluate energy harvesting adaptive algorithms we designed for energy harvesting nodes and links (i.e., the contributions we describe in Chapter 5) with light energy traces we collected (i.e., the contributions we describe in Chapter 3). Additionally, using the testbed, we evaluated a set of heuristic policies for networks of energy harvesting adaptive nodes. To the best of our knowledge, our work is the first attempt to evaluate energy harvesting adaptive policies in a controllable experimental environment. Our evaluations demonstrate a close agreement of the

⁷We provided the prototypes with light inputs corresponding to the light energy recorded over a day in L-1 and L-2 (nearby locations in the same office), compressed to 400 seconds (6.6 minutes) and scaled by 20.0x. The prototypes were equipped with OPVs. In 8 experiments, the variability in total energy harvested was under 3% for all nodes.

Prototype	Sept. 2	2009										Apr. 2
Energy-harvesting aware algorithms	 Node 	Link	÷₹	Single-hop network		£	≯	Single-hop petwor	2		*	
Energy harvesting	 Light sensor 	Crystalline Si solar cells	✦	Amorphous Si solar cells	+	-	→		•	Multihop network	*	Flexible solar cells
+ solar cells Communications	 Sensor 	board	≁	Battery + charge monitoring	≁	Real-time harvested energy monitoring	→	Organic solar cell	*		≯	Flexible batteries
+ Medium Access Control (MAC)	CC100	0 radio				UWB-IR Transceiver UWB-IR MAC	≯		>		*	
Testhed	Flash	g monitoring lights and desk lamps				Manual light control,		Manual light	→ →	t	→	
restord	ι					incundescent namps		control, LEDs		Software-based light control (SBLC)	≁	Compact SBLC
			Phase I	I	Phase II		Phase	III	Phas	e IV	Phase V	Phase VI

Figure 49: A schematic diagram of phase-based EnHANTs prototype and testbed development.



Figure 50: EnHANTs prototypes and testbed, Phases I-V.

experimental results with the simulations, confirm that in many cases simple node and link policies perform well, and demonstrate *experimentally* the tightness of some of the bounds on policy performance we derived in Chapter 5. The evaluations also highlight the need for flow control and topology control policies for networks of energy harvesting nodes.

6.A Phased Prototype and Testbed Development

We have developed the EnHANTs prototypes and testbed in a set of 6 "evolutionary" phases over 4 years, as shown schematically in Fig. 49. At the end of each phase, a fully functional EnHANTs prototypes testbed was showcased at a major conference demonstration session [26,28,57,80,85,119]. The photos of the testbed in Phases I-V are shown in Fig. 50.

The EnHANTs were first prototyped using commercial off-the-shelf sensor network motes [110] and were later iteratively integrated with the custom-designed hardware. The EnHANTs testbed was originally a simple data logger with a bare-bones visualization interface. Later, the testbed was iteratively upgraded to provide real time energy harvesting parameter monitoring and software-based light control functionality.

Below, we present a brief summary of the evolutionary EnHANTs prototypes functionality development:

- Energy harvesting Initially, we designed the EnHANTs prototypes to sense, but not harvest, the available environmental energy (Phase I). Next we integrated commercial solar cells with the prototypes and implemented real-time energy harvesting state monitoring (Phases II and III). We then integrated the custom-designed solar cells (Phase IV). In the latest phase, we integrated the prototypes with *mechanically flexible* solar cells and *mechanically flexible* batteries (Phase VI).
- Ultra-Wideband Impulse-Radio (UWB-IR) wireless communications Prior to the integration of the custom-designed UWB-IR communication modules in Phase III, we substantially modified the mote operating system (which did not support custom transceivers). The integration additionally required the implementation of a custom medium access control module, since the UWB-IR transceiver characteristics differ greatly from the properties of the conventional transceivers (e.g., the clear channel assessment functionality, which is taken "for granted" in conventional transceivers, is not straightforward with UWR-IR).
- Energy-harvesting-aware algorithms The algorithms were first designed and developed for simple single node scenarios, and were later implemented for network scenarios. Following the integration of the UWB-IR transceivers in Phase III, we re-implemented the algorithms to take the UWB-IR characteristics into account.
- Testbed functionality The EnHANTs testbed first consisted of a data logger with a simple visualization interface, which we replaced with a custom-designed real-time monitoring and control system. We additionally developed several light energy control systems, from relatively simple manual setups (Phases III and IV) to a software-based light control system that exposes the prototypes to real-world trace-based light energy conditions (Phase V). Subsequently, we additionally redesigned the software-based light control system to be compact and user-friendly (Phase VI).

Chapter 7

Conclusions

In this thesis, we consider energy harvesting networked nodes. We make contributions to *energy* source characterization and algorithm design for ultra-low-power energy harvesting nodes. In addition, we contribute to node design and prototyping and to project-based engineering education.

We characterized *light energy* and *kinetic (motion) energy* for ultra-low-power energy harvesting nodes. Our characterizations are based on a diverse set of light and motion measurements that we conducted. The characterizations provide important insights into light and motion energy availability and properties (e.g., variability, predictability, influencing factors). The insights are useful for designing energy harvesting nodes and energy harvesting adaptive algorithms. We also formulated and studied *resource allocation problems* for energy harvesting networked nodes. Inspired by the needs of tracking and monitoring Internet of Things applications of networked nodes, we aimed to allocate nodes' varying energy in a uniform way with respect to time. For the deterministic energy profile and stochastic environmental energy models, for node and link scenarios, we formulated optimization problems and introduced algorithms for solving them. We also examined many simple policies, for many of which we provided performance guarantees. Additionally, we designed a new type of ultra-low-power wireless nodes – Energy Harvesting Active Networked Tags (EnHANTs). We prototyped the EnHANTs and designed and developed an EnHANTs prototypes testbed. We used the testbed to evaluate experimentally the energy harvesting adaptive policies we developed with the light energy traces we collected. The design and performance evaluation insights we obtained apply beyond EnHANTs to networks of different energy harvesting nodes. Finally, we also explored new approaches to engaging students in large-scale interdisciplinary research efforts and demonstrated the effectiveness of our approaches.

Bibliography

- M. I. Ali, B. M. Al-Hashimi, J. Recas, and D. Atienza, "Evaluation and design exploration of solar harvested energy prediction algorithm," in *Proc. DATE'10*, Mar. 2010.
- [2] M. A. Antepli, E. Uysal-Biyikoglu, and H. Erkal, "Optimal scheduling on an energy harvesting broadcast channel," in *Proc. IEEE WiOpt'11*, May 2011.
- [3] N. Bambos, "Toward power-sensitive network architectures in wireless communications: concepts, issues, and design aspects," *IEEE Pers. Commun.*, vol. 5, no. 3, pp. 50–59, June 1998.
- [4] S. Bandyopadhyay and A. P. Chandrakasan, "Platform architecture for solar, thermal, and vibration energy combining with MPPT and single inductor," *IEEE J. Solid-State Circuits*, vol. 47, no. 9, pp. 2199–2215, June 2012.
- [5] S. Beeby, M. Tudor, and N. White, "Energy harvesting vibration sources for microsystems applications," *Measurement science and technology*, vol. 17, no. 12, 2006.
- [6] A. Bernat, P. J. Teller, A. Gates, and N. Delgado, "Structuring the student research experience," in *Proc. ACM ITiCSE'00*, July 2000.
- [7] D. Bertsekas and R. Gallager, *Data Networks*, 2nd ed. Prentice-Hall, 1992.
- [8] T. Buren, P. Lukowicz, and G. Troster, "Kinetic energy powered computing: experimental feasibility study," in *Proc. IEEE ISWC'03*, Feb. 2003.
- [9] P. Castiglione, O. Simeone, E. Erkip, and T. Zemen, "Energy-neutral source-channel coding in energy-harvesting wireless sensors," in *Proc. IEEE WiOpt'11*, May 2011.
- [10] S. Chen, P. Sinha, N. Shroff, and C. Joo, "Finite-horizon energy allocation and routing scheme in rechargeable sensor networks," in *Proc. IEEE INFOCOM'11*, Apr. 2011.
- [11] —, "A simple asymptotically optimal energy allocation and routing scheme in rechargeable sensor networks," in *Proc. IEEE INFOCOM'12*, Mar. 2012.
- [12] C.-S. Chow, "Multigrid algorithms and complexity results for discrete-time stochastic control and related fixed-point problems," Ph.D. dissertation, MIT, Cambridge, MA, 1990.
- [13] B. Coleman and M. Lang, "Collaboration across the curriculum: A disciplined approach to developing team skills," in *Proc. ACM SIGCSE'12*, Feb. 2012.
- [14] M. Crepaldi, C. Li, K. Dronson, J. Fernandes, and P. Kinget, "An ultra-low-power interferencerobust IR-UWB transceiver chipset using self-synchronizing OOK modulation," in *Proc. IEEE ISSCC'10*, Feb. 2010.

- [15] —, "An interference robust S-OOK IR-UWB transceiver chipset," IEEE J. Solid-State Circuits, vol. 46, no. 10, pp. 2284–2299, Oct. 2011.
- [16] D. Daly, P. Mercier, M. Bhardwaj, A. Stone, Z. Aldworth, T. Daniel, J. Voldman, J. Hildebrand, and A. Chandrakasan, "A pulsed UWB receiver SoC for insect motion control," *IEEE J. Solid-State Circuits*, vol. 45, no. 1, pp. 153–166, Jan. 2010.
- [17] N. Dang, E. Bozorgzadeh, and N. Venkatasubramanian, "QuARES: Quality-aware data collection in energy harvesting sensor networks," in *IEEE IGCC'11*, June 2011.
- [18] P. Dutta, J. Hui, J. Jeong, S. Kim, C. Sharp, J. Taneja, G. Tolle, K. Whitehouse, and D. Culler, "Trio: Enabling sustainable and scalable outdoor wireless sensor network deployments," in *Proc. IEEE IPSN'06*, Apr. 2006.
- [19] 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.
- [20] 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.
- [21] M. Gorlatova, A. Bernstein, and G. Zussman, "Performance evaluation of resource allocation policies for energy harvesting devices," in *Proc. IEEE WiOpt'11*, May 2011.
- [22] —, "Performance evaluation of resource allocation policies for energy harvesting devices," Columbia University, Tech. Rep. 2011-01-05, Jan. 2011, available at http://enhants.ee. columbia.edu/images/papers/cu_ee_2011-01-05.pdf.
- [23] M. Gorlatova, P. Kinget, I. Kymissis, D. Rubenstein, X. Wang, and G. Zussman, "Energyharvesting active networked tags (EnHANTs) for ubiquitous object networking," *IEEE Wireless Commun.*, vol. 17, no. 6, pp. 18–25, Dec. 2010.
- [24] M. Gorlatova, R. Margolies, J. Sarik, G. Stanje, J. Zhu, B. Vigraham, M. Szczodrak, L. Carloni, P. Kinget, I. Kymissis, and G. Zussman, "Prototyping energy harvesting active networked tags (EnHANTs)," Columbia University, Tech. Rep. 2012-07-27, July 2012, available at http: //enhants.ee.columbia.edu/images/papers/cu-ee-2012-07-27.pdf.
- [25] —, "Prototyping energy harvesting active networked tags (EnHANTs)," in Proc. IEEE INFOCOM'13 mini-conference, Apr. 2013.
- [26] M. Gorlatova, Z. Noorbhaiwala, A. Skolnik, J. Sarik, M. Szczodrak, J. Chen, M. Zapas, L. Carloni, P. Kinget, I. Kymissis, D. Rubenstein, and G. Zussman, "Prototyping energy harvesting active networked tags: Phase II MICA mote-based devices," *ACM MobiCom'10 Demo Session*, Sept. 2010.
- [27] M. Gorlatova, J. Sarik, P. Kinget, I. Kymissis, and G. Zussman, "Project-based learning within a large-scale interdisciplinary research effort," in *Proc. ACM ITiCSE'13*, July 2013.
- [28] M. Gorlatova, T. Sharma, D. Shrestha, E. Xu, J. Chen, A. Skolnik, D. Piao, P. Kinget, I. Kymissis, D. Rubenstein, and G. Zussman, "Prototyping Energy Harvesting Active Networked Tags (EnHANTs) with MICA2 motes," in *Proc. IEEE SECON'10*, June 2010.
- [29] M. Gorlatova, A. Wallwater, and G. Zussman, "Networking low-power energy harvesting devices: Measurements and algorithms," Columbia University, Tech. Rep. 2010-12-15, Dec. 2010, available at http://enhants.ee.columbia.edu/images/papers/CU_EE_2010_12_15.pdf.

- [30] —, "Networking low-power energy harvesting devices: Measurements and algorithms," in *Proc. IEEE INFOCOM'11*, Apr. 2011.
- [31] —, "Networking low-power energy harvesting devices: Measurements and algorithms," To appear in IEEE Trans. Mobile Comput., 2013.
- [32] 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 Mobi-Com*'09, Sept. 2009.
- [33] M. Gorlatova, M. Zapas, E. Xu, M. Bahlke, I. J. Kymissis, and G. Zussman, "CRAWDAD data set columbia/enhants (v. 2011-04-07)," http://crawdad.cs.dartmouth.edu/columbia/enhants, Apr. 2011.
- [34] T. Greening and J. Kay, "Undergraduate research experience in computer science education," in Proc. ACM ITiCSE'02, June 2002.
- [35] 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.
- [36] 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.
- [37] F. S. Hillier and G. J. Lieberman, Introduction to Operations Research, 6th ed. McGraw-Hill, 1995.
- [38] 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.
- [39] H. Huang, G. Merrett, and N. White, "Human-powered inertial energy harvesters: The effect of orientation, location and activity on obtainable power," in *Proc. Eurosensors'11*, Sept. 2011.
- [40] L. Huang and M. Neely, "Utility optimal scheduling in energy harvesting networks," in Proc. ACM MobiHoc'11, May 2011.
- [41] S. Hussmann and D. Jensen, "Crazy car race contest: Multicourse design curricula in embedded system design," *IEEE Trans. Educ.*, vol. 50, no. 1, pp. 61–67, 2007.
- [42] N. Jaggi, K. Kar, and A. Krishnamurthy, "Rechargeable sensor activation under temporally correlated events," in *Proc. IEEE WiOpt'07*, Apr. 2007.
- [43] X. Jiang, J. Polastre, and D. Culler, "Perpetual environmentally powered sensor networks," in Proc. IEEE IPSN'05, Apr. 2005.
- [44] A. Johansen, T. Sorensen, and P. Bonnet, "Service and experiment: Towards a perpetual sensor network testbed without backchannel," in *Proc. IEEE MASS'11*, Oct. 2011.
- [45] C. E. Jones, K. M. Sivalingam, P. Agrawal, and J. C. Chen, "A survey of energy efficient network protocols for wireless networks," ACM/Kluwer Wireless Networks, vol. 7, no. 4, pp. 343–358, 2001.
- [46] 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.

- [47] 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.
- [48] 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.
- [49] M. Khouzani, S. Sarkar, and K. Kar, "Optimal routing and scheduling in multihop wireless renewable energy networks," in *Proc. IEEE ITA'11*, Feb. 2011.
- [50] 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.
- [51] J. Kymissis, C. Kendall, J. Paradiso, and N. Gershenfeld, "Parasitic power harvesting in shoes," in *Proc. IEEE ISWC'98*, 1998.
- [52] C. Lee, J. Su, K. Lin, J. Chang, and G. Lin, "A project-based laboratory for learning embedded system design with industry support," *IEEE Trans. Educ.*, vol. 53, no. 2, pp. 173–181, 2010.
- [53] N. Li, L. Chen, and S. Low, "Optimal demand response based on utility maximization in power networks," in *IEEE Power and Energy Society General Meeting*, July 2011.
- [54] 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.
- [55] 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.
- [56] X. Lu and S. Yang, "Solar energy harvesting for ZigBee electronics," in Proc. SEB'09, Apr. 2009.
- [57] R. Margolies, L. Pena, K. Kim, Y. Kim, M. Wang, M. Gorlatova, J. Sarik, J. Zhu, P. Kinget, I. Kymissis, and G. Zussman, "Demo: An adaptive testbed of energy harvesting active networked tags (EnHANTs) prototypes," in *Proc. IEEE INFOCOM'13*, Apr. 2013.
- [58] P. Martin, Z. Charbiwala, and M. Srivastava, "DoubleDip: Leveraging thermoelectric harvesting for low power monitoring of sporadic water use," in *Proc. ACM SenSys'12*, Apr. 2012.
- [59] V. Mhatre and C. Rosenberg, "Energy and cost optimizations in wireless sensor networks: A survey," in *Performance Evaluation and Planning Methods for the Next Generation Internet*. Kluwer, 2005.
- [60] P. D. Mitcheson, T. C. Green, E. M. Yeatman, and A. S. Holmes, "Architectures for vibrationdriven micropower generators," *IEEE/ASME J. Microelectromech. Syst.*, vol. 13, no. 3, pp. 429–440, 2004.
- [61] P. Mitcheson, E. Yeatman, G. Rao, A. Holmes, and T. Green, "Energy harvesting from human and machine motion for wireless electronic devices," *Proc. of the IEEE*, vol. 96, no. 9, pp. 1457–1486, Sept. 2008.
- [62] 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.

- [63] D. Niyato, E. Hossain, M. Rashid, and V. Bhargava, "Wireless sensor networks with energy harvesting technologies: A game-theoretic approach to optimal energy management," *IEEE Wireless Commun.*, vol. 14, no. 4, pp. 90–96, Aug. 2007.
- [64] D. Noh, L. Wang, Y. Yang, H. Le, and T. Abdelzaher, "Minimum variance energy allocation for a solar-powered sensor system," in *Proc. DCOSS'09*, June 2009.
- [65] D. Noh and T. Abdelzaher, "Efficient flow-control algorithm cooperating with energy allocation scheme for solar-powered WSNs," Wireless Comm. and Mobile Comput., vol. 12, no. 5, pp. 379–392, 2012.
- [66] M. Orendurff, J. Schoen, G. Bernatz, A. Segal, and G. Klute, "How humans walk: Bout duration, steps per bout, and rest duration," *J. Rehabil. Res. Dev.*, vol. 45, no. 7, pp. 1077–89, 2008.
- [67] 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.
- [68] J. Paradiso and T. Starner, "Energy scavenging for mobile and wireless electronics," *IEEE Pervasive Comput.*, vol. 4, no. 1, pp. 18–27, 2005.
- [69] B. Parida, S. Iniyan, and R. Goic, "A review of solar photovoltaic technologies," *Renewable and Sustainable Energy Reviews*, vol. 15, no. 3, pp. 1625–1636, Apr. 2011.
- [70] J. Peckham, P. Stephenson, J. Hervé, R. Hutt, and M. Encarnação, "Increasing student retention in computer science through research programs for undergraduates," ACM SIGCSE Bull., vol. 39, no. 1, pp. 124–128, 2007.
- [71] P. Peumans, A. Yakimov, and S. Forrest, "Small molecular weight organic thin-film photodetectors and solar cells," J. Appl. Phys., vol. 93, p. 3693, 2003.
- [72] J. Polack-Wahl and K. Anewalt, "Learning strategies and undergraduate research," ACM SIGCSE Bull., vol. 38, no. 1, pp. 209–213, 2006.
- [73] M. L. Puterman, Markov Decision Processes: Discrete Stochastic Dynamic Programming. Wiley, 1994.
- [74] 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.
- [75] D. Raicu and J. Furst, "Enhancing undergraduate education: A REU model for interdisciplinary research," ACM SIGCSE Bull., vol. 41, no. 1, pp. 468–472, 2009.
- [76] M. Raju, "Energy harvesting. ULP meets energy harvesting: a game-changing combination for design engineers," Texas Instruments, http://focus.ti.com/lit/wp/slyy018/slyy018.pdf, 2008.
- [77] J. Randall, Designing Indoor Solar Products, 1st ed. Wiley, 2005.
- [78] R. Rao, S. Vrudhula, and D. Rakhmatov, "Battery modeling for energy aware system design," *IEEE Computer*, vol. 36, no. 12, pp. 77–87, 2003.
- [79] S. Russel, The Architecture of Light. Conceptnine, 2008.

- [80] J. Sarik, L. Pena, M. Wang, K. Kim, H. Wang, F. Duque, G. Burrow, R. Margolies, M. Gorlatova, B. Vigraham, P. Kinget, I. Kymissis, and G. Zussman, "Demo: An adaptive testbed of energy harvesting active networked tags (EnHANTs) prototypes," in *IDTechEx Energy Harvesting and Storage USA'12*, Nov. 2012.
- [81] S. Schocken, "Taming complexity in large-scale system projects," in Proc. ACM SIGCSE'12, Feb. 2012.
- [82] N. Sharma, J. Gummeson, D. Irwin, and P. Shenoy, "Cloudy computing: leveraging weather forecasts in energy harvesting sensor systems," in *Proc. IEEE SECON'10*, June 2010.
- [83] V. Sharma, U. Mukherji, V. Joseph, and S. Gupta, "Optimal energy management policies for energy harvesting sensor nodes," *IEEE Trans. Wireless Comm.*, vol. 9, no. 4, pp. 1326–1336, 2010.
- [84] S. Solda, M. Caruso, A. Bevilacqua, A. Gerosa, D. Vogrig, and A. Neviani, "A 5 Mb/s UWB-IR transceiver front-end for wireless sensor networks in 0.13 CMOS," *IEEE J. Solid-State Circuits*, vol. 46, no. 7, pp. 1636–1647, July 2011.
- [85] 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." in *Proc. ACM SenSys'11*, Nov. 2011.
- [86] T. Starner, "Human-powered wearable computing," *IBM systems Journal*, vol. 35, no. 3, pp. 618–629, 1996.
- [87] M. Szczodrak and L. Carloni, "Demo: A complete framework for programming event-driven, self-reconfigurable low power wireless networks," in *Proc. ACM SenSys'11*, Nov. 2011.
- [88] J. Taneja, J. Jeong, and D. Culler, "Design, modeling, and capacity planning for micro-solar power sensor networks," in *Proc. IEEE IPSN'08*, Apr. 2008.
- [89] K. Tutuncuoglu and A. Yener, "Optimum transmission policies for battery limited energy harvesting nodes," *IEEE Trans. Wireless Commun.*, vol. 11, no. 3, pp. 1180–1189, Mar. 2012.
- [90] "Cymbet EnerChips," http://www.cymbet.com/content/products.asp.
- [91] "EnOcean Energy Harvesting Wireless Sensor Solutions and Networks," http://www.enocean. com/.
- [92] "GreenPeak Technologies," http://www.greenpeak.com/.
- [93] "Infinite power solutions evaluation kits," http://www.infinitepowersolutions.com/products/ evalkits.html.
- [94] "Intel Lab data," http://db.csail.mit.edu/labdata/labdata.html.
- [95] "Measurement and Instrumentation Data Center, National Renewable Energy Laboratory (NREL), US DOE," http://www.nrel.gov/midc/.
- [96] "nPower Personal Energy Generator," http://www.npowerpeg.com/.
- [97] "Powercast development kits," http://www.powercastco.com/products/development-kits/.

- [98] "Preparing the Plexcore PV 1000 Ink System," http://www.sigmaaldrich.com/ technical-documents/protocols/materials-science/preparing-the-plexcore.html.
- [99] "TAOS TSL230rd programmable light-to-frequency converter," http://www.taosinc.com/.
- [100] "Texas Instruments MSP430 Solar Energy Harvesting Development Tool," http://focus.ti. com/docs/toolsw/folders/print/ez430-rf2500-seh.html.
- [101] "Weather Underground Weather History and Data Archive," http://www.wunderground.com/ history/.
- [102] 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.
- [103] T. Von Buren, P. Mitcheson, T. Green, E. Yeatman, A. Holmes, and G. Troster, "Optimization of inertial micropower generators for human walking motion," *IEEE Sensors J.*, vol. 6, no. 1, pp. 28–38, 2006.
- [104] R. Vullers, R. Van Schaijk, I. Doms, C. Van Hoof, and R. Mertens, "Micropower energy harvesting," *Solid-State Electronics*, vol. 53, no. 7, pp. 684–693, Apr. 2009.
- [105] Z. Wang, A. Tajer, and X. Wang, "Communication of energy harvesting tags," *IEEE Trans. Commun.*, vol. 60, no. 4, pp. 1159–1166, Apr. 2012.
- [106] B. Warneke, M. Last, B. Liebowitz, and K. Pister, "Smart dust: communicating with a cubicmillimeter computer," *IEEE Computer*, vol. 34, no. 1, pp. 44–51, Jan 2001.
- [107] E. Wenderholm, "Challenges and the elements of success in undergraduate research," ACM SIGCSE Bull., vol. 36, no. 4, pp. 73–75, 2004.
- [108] 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.
- [109] Y. Wu, Z. Mao, S. Fahmy, and N. Shroff, "Constructing maximum-lifetime data gathering forests in sensor networks," *IEEE/ACM Trans. Netw.*, vol. 18, no. 5, pp. 1571–1584, 2010.
- [110] N. Xu, "A survey of sensor network applications," *IEEE Commun. Mag.*, vol. 40, no. 8, pp. 102–114, 2002.
- [111] Y. Xue and L. Jin, "A naturalistic 3D acceleration-based activity dataset and benchmark evaluations," in *Proc. IEEE SMC'10*, Oct. 2010.
- [112] J. Yang, O. Ozel, and S. Ulukus, "Broadcasting with a battery limited energy harvesting rechargeable transmitter," in *Proc. IEEE WiOpt'11*, May 2011.
- [113] J. Yang and S. Ulukus, "Optimal packet scheduling in an energy harvesting communication system," *IEEE Trans. Commun.*, vol. 60, no. 1, pp. 220–230, Jan. 2012.
- [114] 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.
- [115] L. Yerva, A. Bansal, B. Campbell, T. Schmid, and P. Dutta, "Grafting energy-harvesting leaves onto the sensornet tree," in *Proc. IEEE IPSN'12*, Apr. 2012.

- [116] J. Yun, S. Patel, M. Reynolds, and G. Abowd, "Design and performance of an optimal inertial power harvester for human-powered devices," *IEEE Trans. Mobile Comput.*, vol. 10, no. 5, pp. 669–683, May 2011.
- [117] M. Zafer and E. Modiano, "A calculus approach to energy-efficient data transmission with quality-of-service constraints," *IEEE/ACM Trans. Netw.*, vol. 17, no. 3, pp. 898–911, 2009.
- [118] D. Zhu and S. Beeby, Energy Harvesting Systems: Principles, Modeling, and Applications, 1st ed. Springer, 2011.
- [119] 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," in *Proc. ACM MobiSys'11 Demo Session*, June 2011.
- [120] 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.

Appendix A

Project-based Learning within the EnHANTs Project

In this appendix, we describe our experiences with engaging a large and diverse group of students in *project-based learning* within the EnHANTs prototype and testbed design and development "*umbrella project*". While project-based learning is actively used to help students build professional skills [41,52], typically it is only applied to small teams and small efforts. We, on the other hand, involved a diverse population of over 50 students on over 100 semester-long projects in our effort. To the best of our knowledge, our experience with organizing multiple student projects to contribute to a large-scale effort is unique. The results are based on joint work with J. Sarik. R. Margolies contributed to many of the organizational approaches discussed in this appendix.

We first describe the student projects (Section A.1) and our approaches to organizing them (Section A.2). We then summarize some of the lessons learned (Section A.3) and present the evaluation results (Section A.4).

We previously presented the results that appear in this appendix in [27].



Figure 51: Student projects completed in different semesters.



Figure 52: Students involved in the EnHANTs project.

A.1 Student Projects

Under the EnHANTs umbrella project, over 11 semesters, 52 students completed 115 projects. The number of student projects completed in each semester is shown in Fig. 51. Student demographics are presented in Fig. 52 and 53. Out of 51 students, 6% were high school students we engaged via Harlem Children Society. 31% were undergraduates, 50% were M.S. students (of which all but one were in non-thesis terminal M.S. programs), and 13% were Ph.D. students. 75% of students were enrolled in academic programs in Columbia University, while the other 25% were visiting students, i.e., Research Experience for Undergraduate (REU) students, students from local colleges without advanced research facilities, or visiting international students. Out of 115 student projects, 51% were full-time projects (summer research internships, REU projects, M.S. thesis research semesters). The other 49% were semester-long research project courses to which students typically dedicated 8-15 hours per week. 70% of the projects were completed by male students and 30% by female students. The main focus areas for the student projects were networking (52%), circuits and systems (25%), electronics and applied physics (15%), and operating systems (8%).

The student projects within the EnHANTs project were *collaborative and multidisciplinary*. A project typically focused on one disciplinary area (e.g., algorithm design, operating systems development, solar cell design), but required interaction with at least two other areas. These projects



Figure 53: Students semesters.

challenged students by requiring them to gain understanding of concepts outside of their comfort zone. Additionally, students improved their communication and teamwork skills because the projects required them to independently and proactively seek out relevant expertise throughout the research groups involved in the EnHANTs project. Finally, the projects exposed the students to all aspects of networking, from the physical-layer pulses generated by the UWB-IR transceivers to the adaptive flow control and routing protocols. Students thus gained an *in-depth fundamental understanding of networking concepts*. A few representative student projects are described below.

- Real-time monitoring and control system, completed by an undergraduate Computer Science student A. Skolnik: The student developed a Java-based system to monitor and control the EnHANTs prototypes. The project involved designing the necessary data structures to enable communication between the prototypes and the computer running the monitoring system. The student designed the system to support both a text-based interface and a "visual demo" interface that shows the activity of the prototypes in an easy-to-understand way. This project, implemented using TinyOS and Java, required knowledge of sensor networking, wireline communication, and software design. The student extensively interacted with students who were modifying the prototype operating system and developing energy-harvesting-aware algorithms.
- EnHANTs prototype UWB-IR communication module, completed by a M.S. Electrical Engineering student J. Zhu: The student developed and tested the UWB-IR communication module. The student integrated a custom-designed UWB-IR transmitter and receiver chipset onto a single printed circuit board and programmed a complex programmable logic device to perform data serialization and deserialization, preamble detection, and byte synchronization. The student also developed a UWB-IR radio driver using TinyOS. Primarily focused on circuit design, this project

required the student to develop expertise in networking, operating systems, and software design.

• Energy-harvesting-adaptive EnHANTs network, completed by an undergraduate Computer Engineering student D. Roggensinger: The student implemented network layer protocols that handled the EnHANTs packet routing. The student first tested the network functionality using commercial transceivers, and then extensively evaluated network's performance with custom UWB-IR transceivers. The student also implemented energy harvesting adaptive network layer algorithms, which were adapting packet routing paths based on the environmental energy availability. The student extensively tested these algorithms with the custom energy harvesting modules. While primarily focused on networking, the project required the student to gain an in-depth knowledge of the UWR-IR communications and energy harvesting.

A.2 Project Organization

Organizing multiple student projects to contribute to a large-scale effort is challenging. We present some of our approaches to organizing projects, motivating students, and facilitating learning. **Real-world system integration deadlines:** EnHANTs prototype and testbed design, development, and integration proceeded in a series of phases (see Fig. 49 and 50). At the end of each phase, the fully integrated EnHANTs prototype and testbed were presented at a major conference [26, 28, 57, 80, 85, 119]. We used the conference timelines as the "real-world" deadlines for the integration of different student projects. The benefits of this approach are multi-fold. First of all, it *motivates students*. Providing short-term deadlines for student projects, rather than abstract longterm goals, energizes and motivates the students. Students are additionally motivated by seeing their work integrated with the work of others, used in a conference presentation, and subsequently extended. It also *encourages cross-disciplinary collaboration*, as under short-term system integration deadlines, the students work, individually and jointly, to quickly solve problems as they arise. Finally, it *reduces the impact of unsuccessful projects*. By constantly updating the software and the hardware components throughout the system integration deadlines, we restrict the negative impact of the projects that are technologically flawed.

Frequent cross-group meetings: We conducted regular (weekly or bi-weekly) meetings where students presented their work to the faculty and students from the different research groups. This challenged students to present their work so that it could be understood by people with different backgrounds. Additionally, students reported that observing how faculty members solved problems during these meetings improved their own problem solving skills.

Ph.D. student mentorship: The faculty members involved in the EnHANTs project were heavily engaged in the student projects. However, faculty members delegated many of the day-to-day student supervision tasks to their Ph.D. students. The Ph.D. students provided technical support and guidance to the students, tested and verified student projects before integration with EnHANT prototypes and testbed, and ensured continuity among the different student projects. While somewhat time-consuming, these tasks provided the Ph.D. students with important opportunities to demonstrate and improve their mentorship, leadership, and project management skills.

Frequent system demonstrations: Functional ("live") EnHANT prototypes and testbed were frequently demonstrated in different on-site and off-site presentations.¹ Frequent demonstrations, particularly those conducted off-site, encouraged students to design and develop robust software, hardware, and algorithms and to extensively verify and test their work. This improved students' technical skills, and provides them with an understanding of the quality standards required from technology in "real-world" applications. Additionally, the testbed demonstrations gave students opportunities to present their work to vastly different audiences.

A.3 Lessons Learned

Throughout the 4-year course of the EnHANT project we learned many important lessons. Our experiences highlight and reinforce the need to foster opportunities for *close and continuous cross-group interactions*.

The students work in different labs, focus on different disciplines, and have different technical skills, priorities, work styles, and expectations. Early on we discovered that the gaps between the knowledge of the students with different expertise areas are much wider than anticipated. For example, Electrical Engineering students are oftentimes unfamiliar with good software development

¹Inspired by *agile software development practices*, we ensured that a version of the EnHANT prototypes and testbed was always ready to be demonstrated. We did not integrate new software or hardware without extensive testing and design for backward compatibility. This further reduced the impact of unsuccessful student projects.

practices, while many Computer Science students may not understand how to properly handle experimental electronics. Additionally, many students that are not majoring in Electrical Engineering do not understand the concepts of frequency-domain signal processing that are essential to the understanding of wireless networking. These knowledge gaps often lead to both technological and interpersonal issues. Cross-group problem solving requires students to trust each other's expertise, but these gaps in knowledge can make the trust difficult to establish. When working with students we highlight that such gaps are normal and should be treated as a learning opportunity. Ultimately these issues can only be addressed by *establishing, maintaining, and nurturing the connections between the groups and between the different students*. Specifically, we have learned the importance of the following:

- *Extensive use of collaborative tools*: As the project has progressed, we have expanded our use of collaborative tools. The EnHANTs project has an internal wiki and an external website that are kept up to date with shared technical information and documentation. We use less formal Google docs to keep the students "on the same page" with regards to project timelines and tasks.
- Carefully defined interfaces between student projects: Some of the most challenging problems arise when a student interfaces his or her project with another project. The difficulty of solving these interface problems can in certain cases lead to interpersonal tensions. Designing the interfaces between different technologies (e.g., solar cells and the energy harvesting module, UWB-IR communication module and the control module) has been a challenging task that often required faculty members' involvement.
- Formalized knowledge transfer process: A large-scale, long-term project necessitates knowledge transfer between the students. In our experience, while knowledge transfer needs to be carefully monitored and emphasized, most students embrace it when they see first-hand that the documentation they create is used by their peers. Similarly, most students embrace the opportunity to introduce peers to their work and to teach them.
- Showcase of individual student contributions: With many student projects integrated into the prototypes, the contributions of some students may not be as visible as the contributions of others. To address this, we conduct workshops where the students present their projects individually. We also separately showcase each student project on the EnHANTs project website.



Figure 54: Project survey results.

A.4 Experiences and Feedback

In October 2012 we conducted a survey amongst all 45 high school, undergraduate, and M.S. students that had participated in the project. The survey contained multiple-choice questions and optional open-ended questions. The survey response rate was 75.5%. In the survey's optional open-ended questions, students shared many observations, comments, and suggestions about the EnHANTs project organization. Fig. 54 and 55 show some of the results.

Overall, the students' experiences were overwhelmingly positive. Over 90% of the students believed their project experience to be rewarding and enriching. Over 85% of the students indicated that working on this project improved their ability to function on multidisciplinary teams and to communicate effectively, made them a better computer scientist or a better engineer, and was a valuable part of their education. 70% of the students indicated that working on the project improved their ability to function on multidisciplinary teams *more than any other activity*. Additionally, 50% of the students indicated that working on the project improved their ability to communicate effectively more than any other activity, and over 40% of the students indicated that this project increased their knowledge of computer networking more than any other activity.

The students' impressions of the project organization features provided additional insights into the features' effectiveness:

• Multidisciplinary projects: Most students enjoyed the multidisciplinary nature of their projects. When specifying what they liked most about the project, over 50% of the students commented on one of its multidisciplinary aspects. One student enjoyed her project being "about both hardware



Figure 55: Project survey results: EnHANTs projects compared with other activities.

and software", and said it was "innovative and challenging to integrate many different aspects in one". One student's favorite thing about his project was the "integration of my work with other parts of the system – felt like a cohesive project that mattered more".

- **Ph.D. student mentorship:** The majority of students appreciated the support provided by their Ph.D. student mentors. Over 80% of the students said that their mentor was approachable and accessible, and provided appropriate guidance.
- Frequent cross-group meetings: Most students appreciated the opportunities for problemsolving and work presentations provided by the regular cross-group meetings. One student noted that "the meetings are an excellent way for putting everything in the big picture." Yet several students also commented that the meetings were unnecessarily long, and suggested that a better meeting structure should be considered for the future projects. To improve the quality of the meeting presentations we encourage students to discuss their presentation with their Ph.D. mentors. We are also considering joint presentations for students from the same research group.
- Frequent system demonstrations: Over 95% of the students indicated that presenting their work was a rewarding and enriching experience. Several students specifically mentioned the presentation skills amongst the skills they acquired or improved while working on the EnHANTs project. One student noted that "the opportunity to present to others was invaluable. Plus it was a lot of fun!"

The majority of the negative feedback focused on insufficient knowledge transfer, and the need for further facilitation of cross-group communications. Several students commented on the insufficient technical introduction to their project. One student stated that *"In the beginning I felt I didn't have* enough support to ask very basic things", and another student noted that "a lot of work goes to waste if you are unable to successfully pass it on to the next person". Students noted that "getting everyone on same page was difficult at times", and said that "not being able to know exactly what others are doing" was an impediment to achieving some of their project goals. Based on this feedback, we have increased our efforts to ensure that students create high-quality, up-to-date documentation, and have been additionally encouraging the students to independently collaborate with each other.

As of April 2013, 55% of the students have graduated (the other 45% are continuing their studies). Of the students who have already graduated, 30% continued to higher-level academic programs. Many students have been accepted to Ph.D. programs in leading universities such as University of Illinois at Urbana Champaign (UIUC), Princeton, Harvard, and Carnegie Mellon. The other 70% of the graduates joined different technology companies, including Microsoft, OPNET, and Oracle. Several students have indicated that working on the EnHANTs project prepared them for some of the challenges they face in their careers. For example, one student noted that the "experience presenting my work has been really helpful in my current job profile", and another highlighted that "being held accountable for deadlines and project completeness helped prepare me for work environment".

A.5 Conclusions

While the modern computing landscape increasingly requires large-scale system engineering skills, such skills are rarely acquired in a typical engineering program. To address this, over the last 4 years, we have been engaging a diverse group of students in research projects within a large-scale interdisciplinary EnHANTs project. As of April 2013, 115 projects have been completed within the EnHANTs "umbrella" project. The projects challenge students' knowledge and organizational and communication skills. Some of the approaches we have taken to facilitate student learning are the "real-world" system development constraints and regular cross-group meetings. Students find the projects rewarding and gain valuable skills. Of the students who completed a survey we developed to evaluate student learning, over 90% indicated the project was rewarding and enriching, and 70% indicated that working on this project improved their ability to function on multidisciplinary teams more than any other activity in their academic career.

Our experience demonstrates the feasibility of engaging diverse groups of students on large-scale

interdisciplinary research efforts, sheds light on some potential pitfalls of such efforts (e.g., inadequate cross-group communication and knowledge transfer), and suggests best practices to overcome these challenges.