Sensors with Lasers: Building a WSN Power Grid

Naveed Anwar Bhatti, Affan A. Syed
Systems and Network Lab (SysNet)
National University of Computer and Emerging Sciences
Islamabad, Pakistan
Email: first.last@sysnet.org.pk

Muhammad Hamad Alizai
University of Engineering and Technology
Peshawar, Pakistan
Email: hamad.alizai@nwfpuet.edu.pk

***Abstract*—We present here a first practical energy distribution architecture that allows us to decouple energy supply from sensing activities in WSN. Such a separation of responsibilities enables us to utilize abundant energy sources distant from the sensing location, allowing unrestricted lifetime and resolving unequal energy consumption in WSN. We demonstrate energy transfer for practical decoupling using low-cost and -footprint, laser μ-power beaming that powers current WSN platforms at 100m of range. We design and implement LAMP: a tiered architecture to manage energy supply to both mesh and clustered WSN deployments using an energy distribution protocol. We evaluate our system to show that, for an additional cost of $29 per mote, LAMP can support perpetual mesh functionality for up to 40 sensors or 120 nodes in clustered operation.**

Keywords*—*Wireless sensor networks, lasers, power-beaming, Wireless energy transference, Capacitors

(a) A typical Body sensor network [23] (b) ETH's Deployment on the Matterhorn [4] (c) ISI's deployment in Costa Rican forest [25]

Fig. 1. Existing WSN deployments: The potential for wireless energy transfer is restricted due to coupled energy and sensing planes.

I. Introduction

It is revealing to observe that the vast majority of research in the WSN community has focused on efficient energy management, albeit the principle motivation of WSN is to provide in-situ sensing capability. This research-bias is due to a fundamental coupling of the energy sources and their distribution (*energy plane*), with the spatio-temporal sensing requirements (*sensing plane*) of an application. This coupling, in turn, is necessitated by the axiom that availability of energy (wall-socket or ambient) is independent of the sensing location.

There are two limitations arising from this coupling of a limited energy resource to a sensing location. First, it precludes the use of an abundant and cheap energy source physically distant from the sensor. A resulting tradeoff occurs between sensing fidelity and greater lifetime achieved by using the large potpourri of energy management techniques developed by the research community [13], [27], [31], [34]. Figure 1(a) shows an example body sensor network where a wireless energy solution can simply use wall power to indefinitely run all sensors. While wireless energy might be inefficient, any cost of grid power wasted when weighed against the benefit of removing lifetime related constraints on sensing or communication is, in our opinion, quite acceptable.

A second limitation, due to varying energy consumption across a WSN, results in non-uniform energy resource depletion. Such a scenario can result in a non-functional WSN even when the global (across the deployment area) energy available is sufficient for its operation. While harvesting energy at sensor node mitigates some of these limitations, the independence of ambient energy from sensor placement still results in a similar conundrum. Figure 1(b) and 1(c) show two such real WSN deployments by ETH-Zurich at Matterhorn and USC/ISI in a Costa Rican rainforest. In both deployments we see excess energy in areas with sunlight (one side of a ridge, and the forest canopy) but are forced to use batteries as *all* locations are not equally provisioned. As a result, despite using state-of-the-art energy efficient mechanisms for extending life times, both deployments need frequent, and difficult, battery replacement.

We propose to *decouple* energy and sensing planes in WSN. This is enabled by the wireless energy research [14], [18] that removes the requirement for the energy source to be on or near the sensor node. By decoupling energy and sensing planes we can, (i) treat energy as an independent, deployment-wide shareable resource, (ii) create an *abstraction* of the energy plane that the sensing plane can use, for example, to request distribution of energy to satisfy sensing requirements, and, (iii) alongside existing energy management techniques, allow energy generation and distribution to evolve independent of the sensing requirements of an application.

We first identify four requirements for a wireless technology to practically enable decoupling in WSN. Any technology must: be wireless with small footprint, have low cost, provide energy at radio comparable range, and deliver enough energy to support existing mote class devices such as TelosB. To this end, we build the first practical, *laser-based* μ-power beaming mechanism that meets these requirements by recharging a depleting node at an order of magnitude higher scale (mWs vs. μWs), that too at an order of magnitude greater range (100m vs. 2-3m) than existing mechanisms (**Section III**).

We then design LAMP: an architecture that decouples energy and sensing planes in WSN using laser μ-power beaming (**Section IV**). LAMP is based on a tiered energy plane that distributes energy from power unconstrained (e.g., via wall-sockets) *master* nodes to *leaves* that perform in-situ sensing. Thus, LAMP enables the use of an abundant and cheap energy source physically distant from the sensor and schedule energy supply to match unequal consumption. We design LAMP to support *any* implementation of the sensing plane: from full mesh operation with homogenous sensors, to the tiered operation with specialized routing and sensing nodes commonly used for practical WSN deployments [9], [24].

We then build a prototype of our energy distribution architecture, delineating the practical limits of energy demand it can support for *existing* sensing architectures. We answer questions regarding its accuracy, scalability, and supported data rate (**Section V**). We thus identify the need for laser recalibration at leaf distances greater than 29.5m, and a minimum spacing of 0.8m between leaves for power beaming at 100m. We show that our implementation can support TelosB class of motes with 7.4% duty cycle at distances in excess of 100m. Furthermore, our system can perpetually support up to 120 leaves per master in clustered mode (only master-leaf communication) and up to 40 leaves in mesh mode (unrestricted intra-leaf communication).

Our work thus has the following three novel contributions:

- we identify essential constraints for practical energy transfer in WSN and build a laser power beaming solution that meets them,
- we design the first energy distribution architecture providing an energy plane abstraction to any sensing application, and
- we implement a prototype of this architecture and thoroughly evaluate it to understand its limitations and scope.

II. Why Decouple Energy distribution from Sensing in WSNs?

We now elaborate on the concept of energy and sensing plane abstraction for a general WSN deployment (Figure 2). We use these abstractions in our paper to later understand the LAMP architecture.

The energy plane (E-plane) represents the distribution of energy sources, such as batteries, wall-sockets, or ambient energy. The sensing plane (S-plane) represents the distribution of application-specific sensory information. Ideally, a WSN application wants to optimize deployment — in terms of cost, lifetime and fidelity — for data sensing with just the S-plane information. However, these two planes do not overlap in most deployments, i.e., an energy rich area may not be a sensing point of interest[1]. This is why current WSN deployments couple these planes by adding energy sources (using batteries) at the sensing location.

[1]energy harvesting is beneficial only in areas where these planes serendipitously overlap

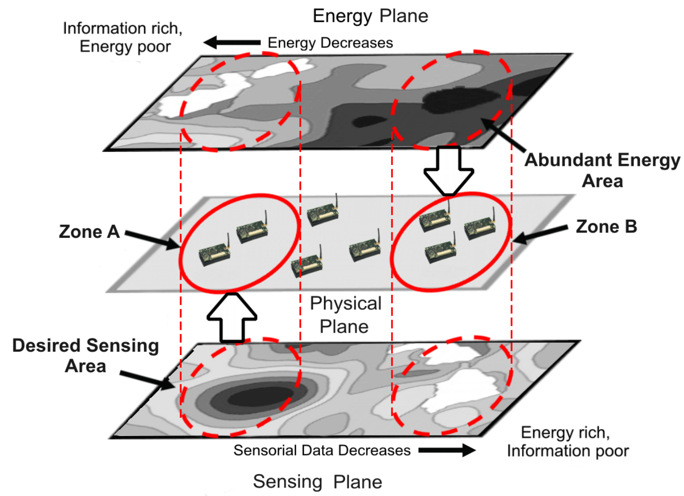

Fig. 2. Decoupled energy and sensing planes in WSN. These planes are spatially uncorrelated in most deployments.

This coupling fundamentally limits energy at sensor nodes due to three application-specific constraints. First, in-situ sensing must be done where the application requires it, but these locations are spatially uncorrelated with close by and abundant energy sources (like sunlight or wall socket). Second, both sensor nodes and their harvesting mechanisms have size limitations to minimize invasiveness of the WSN application (e.g., surveillance or habitat-monitoring). Finally, post-deployment energy consumption varies for nodes with time and usage, causing an unequal distribution of energy at each node.

If we decouple these two planes, we can design a separate architecture that abstracts the transport of energy between locations in the E-plane. We can now, with this abstraction, remove the above limitations in one of two ways: either use it to balance unequal energy distribution in WSN by taking from energy-rich nodes and giving to energy-poor nodes; i.e. a modified Robin-Hood argument. Alternatively, we no longer compromise on the functionality or fidelity of a sensing application just to increase lifetime as energy can now be provided, within reasonable limits, on-demand from a location with abundant energy (like a wall-socket).

A pertinent analogy, and a motivation to decouple energy generation from its use, is the power grid which implements such decoupling, thereby allowing both consumer applications (e.g. home-appliances) and energy distribution technology to evolve independently, and more efficiently.

III. Energy Transfer for WSN using Laser μ-power Beaming

Several recent works propose interesting mechanisms for energy transfer within the context of WSN deployments [6], [17], [18], [26], [33], [37], [38]. However, we believe that each of these mechanisms lacks in some critical aspect, which makes them unsuitable for the majority of practical WSN scenarios. We first identify these minimum requirements for a WSN-specific energy transfer mechanism and then evaluate laser-based μ-power beaming as a suitable option.

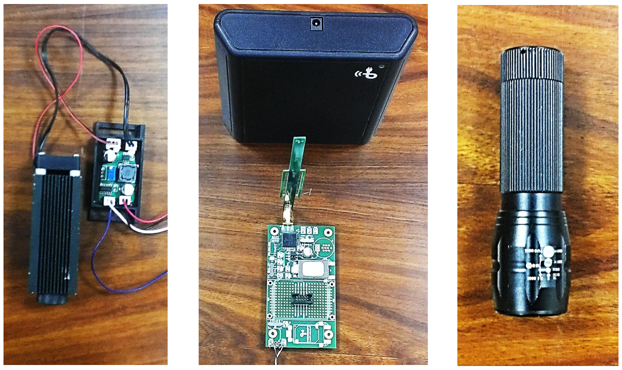

(a) Laser module (0.8W) (b) PowerCast RF module (3W) (c) Cree LED torch (3W)

Fig. 3. Wireless energy technologies.

A. *Practical Constraints for Energy Transfer in WSN*

At the minimum, a practical solution for energy-transfer should not violate the fundamental goals of cheap, in-situ, collaborative, and long lived sensing for WSNs. We thus believe that a WSN-specific energy transfer mechanism should have the following properties: ❶ It should have **small size** and be **wireless** to allow minimally invasive deployments. ❷ It should be a **low-cost** solution (sub \$100), ❸ **providing sufficient energy** (in 10s of mW) to power typical WSN class devices for sensing and communication purposes. Finally, ❹ its **range of energy transfer** should be at least comparable to typical communication range (10s of m; direct or by routing) of WSN nodes to allow reasonable density of deployment.

We acknowledge that there can be other constraints that are specific to an application, like safety, line-of-sight, and directionality. However, the above are minimum requirements for integration with *any* practical WSN application.

B. *Sensors with Lasers*

We propose using laser μ-power beaming as an appropriate mechanism for WSN as it fulfills the above requirements. Our choice of laser μ-power beaming is inspired by higher-density power beaming research for satellites and UAVs [5], [29]. However, in order to meet the cost and size constraints, we radically reduce the scope of our power beaming solution.

We next discuss our setup to evaluate the range and energy transfer capability of such a system. We also experimentally compare our approach with two other approaches for energy transfer proposed earlier: Light-based [18], [26], and RF-based [6], [17].

1) Evaluation Setup: We use three different mechanisms for energy transmission in the context of WSNs (Figure 3). For RF power transfer we use Powercast's module (same as in [6]) that uses a 3W transmitter [22]. We use a 3W Cree Q3 LED torch for energy transfer using light [1]. Finally, for our laser power beaming we use a 808nm 0.8W near-infrared laser module with adjustable beam width [32]. This laser has a driver circuitry to adjust its output power (for safety concerns).

We use a mono-crystalline, high-efficiency solar cell [12] to collect radiant energy transmitted by both laser and LED. This solar cell has high efficiency ($\approx 20\%$) with good response for frequencies of both visible and laser light. Also with its small size ($89 \times 55 \times 2\ mm$) it can be easily attached to any sensor node. For receiving RF power we use Powercast's P2110 IC whose output is regulated by their evaluation board [22]. Note that due to their small size and wireless mode, all three mechanisms meet requirement ❶.

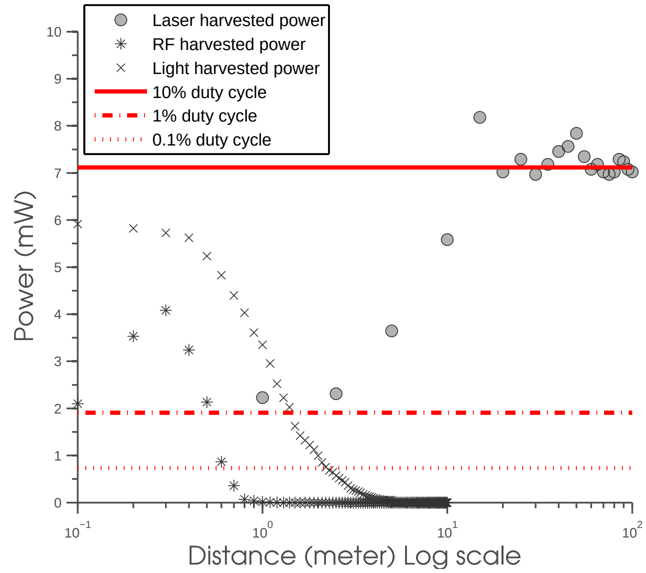

Fig. 4. Harvested power at different ranges (log scale). Also shown are power requirement for TelosB operation at different duty-cycles.

We use a simple setup with the transmitter and receivers perfectly aligned to measure the amount of energy received at different range. For each experiment, we optimally tweak the load at the receiver for fair comparison, as the actual amount of energy harvested depends on the load. Figure 4 shows the result of all the three mechanisms, which we next use to justify the selection of laser power beaming.

2) Range of Power Transfer: The first conclusion we can draw from Figure 4 is that laser power beaming is the only mechanism that can provide consistently good energy transfer at the 80-100m range expected of current WSN nodes (requirement ❹). Thus, while both the RF and light-based transmitter use greater power (3W), they effectively are unable to transfer energy beyond a couple of meters. Laser μ-power beaming, on the other hand, consistently provides around 7mW of power all through 10-100m of range. We observe a decrease in laser power transfer at less than 10m due to the laser beam not covering the entire solar panel at these short distances. This limitation, however, is technological; optics in front of the laser can be used to spread the laser beam and cover the entire panel.

3) Supported Class of WSN Applications: We now evaluate the class of devices or application that each transfer mechanism can support, at their *best operating point* where we normalize the power delivery different ranges. For this purpose we use a simple method where Figure 4 shows three average power requirements for a telosB mote-class device under different duty cycles. These lines can either be viewed to signify different application requirements or a different class of devices [10], [11].

TABLE I. COMPARING WIRELESS ENERGY SCHEMES.

	Laser	Light	RF
Cost (Tx/Rx/Total)	$ 39/28/67	$3/28/**31**	$236/185/421
Range	**+100m**	5-6m	1-2m
Power ([range,power] normalized)	**7mW**	4mW	2-3mW

We again observe that only laser μ-power beaming can support typical WSN application and devices (requirement ❷). We, therefore, do not require a tradeoff for lower capability platforms (like CRFID or EnHANTs [10]) which constrain the WSN applications we can support. Thus, all current WSN applications can simply focus on meeting the sensing requirements and request energy from our energy distribution system. Furthermore, laser μ-power beaming can support these devices even at 100m as its coherent nature results in little loss in power transfer over longer range.

4) Cost and Conclusion: We now evaluate the cost of deploying these solutions on a single pair of devices (to transmit and receive power). Table I shows the cost comparison for deploying the three technologies. We note that laser μ-power beaming is slightly more expensive than simple light based energy transfer. However, this extra cost becomes negligible when amortized over the 120 nodes we can support (see Section V-D).

We thus conclude that laser μ-power beaming is the most suitable energy transfer mechanism among those available that satisfy requirement ❶.

IV. LAMP: DECOUPLING ENERGY AND SENSING PLANES

We now present LAMP, a prototype implementation that allows us to decouple the E-plane, using our laser μ-power beaming mechanism, and transport energy to meet WSN requirements in the S-plane (cf. Section II). We first present an overview of LAMP architecture. Next, we describe the different components involved in the LAMP operation. We then present the LAMP control protocol for energy distribution. We conclude by highlighting LAMP support for different WSN architectures (in the S-plane), thus removing lifetime constraints on a wide range of existing and proposed WSN applications.

A. LAMP Architecture

Once we decouple the energy and sensing planes, an architectural question arises about the energy distribution network. A *homogenous* architecture assumes nodes to be equally capable of sensing, networking, harvesting, and energy transfer. However, questions about practicality of such an architecture, applicable even more for energy transport, have repeatedly raised several ties [8], [16], [20], [24]

A *heterogenous* energy distribution architecture is, in our opinion, more promising. Such a tiered architecture clearly separates nodes responsible for generating energy from the nodes that consume energy in the E-plane and sense in the S-plane. With the help of wireless energy transfer between these nodes, we can then support the desired interaction between E- and S-plane abstractions to satisfy application sensing requirements.

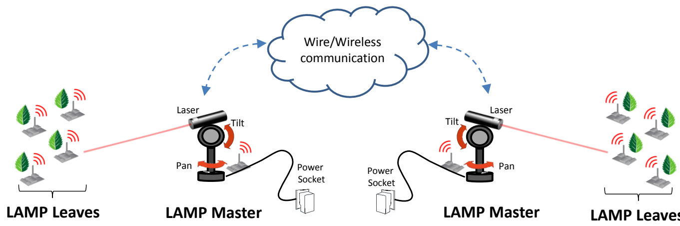

Fig. 5. LAMP Architecture.

1) Tiered Energy Plane: LAMP is thus based on a tiered architecture to achieve the decoupling of the E- and S-planes. As shown in Figure 5, the LAMP master is responsible for energy generation and its distribution to the leaves in the E plane. A master is unconstrained; thus it is engineered with an abundant supply of energy. A LAMP leaf is responsible for receiving energy in the E-plane and sensing and reporting, at the location of interest, in the S-plane. The use of laser power beaming, because of its longer energy-transfer range, allows a relatively unconstrained, spatial deployment of leaves. The master periodically recharges and samples each assigned leaf to achieve the desired goals of a deployment.

We emphasize that this tiering is employed to assign a set of leaves to each master in the E-plane. It does not restrict leaves from communicating with each other or with a sink node, possibly over multiple leaf-hops in the S-plane. We defer a detailed discussion on possible communication patterns in LAMP to Section IV-D

2) Discussion: This architectural separation of node duties in LAMP is not new: it is similar to other tiered architectures in the S plane, such as TENET [20], consisting of constrained *motes* in the lower tier and relatively unconstrained *masters* in the upper tier. TENET introduces this architectural separation to increase network capacity (i.e., using masters) while constraining motes to only minimally process locally-generated data. While promoting TENET motives, LAMP [2] assigns masters an additional role, in the E plane, of generating energy and disseminating it to a set of leaves placed in-situ.

The use of this tiered architecture for LAMP is intentional: it has been repeatedly advocated by WSN pioneers both in the past [7] and of late [24]. Hence, our choice is based on this strong observation that the future WSN deployments will be tiered [8], [16], [19], [20], [24] especially with the recent surge of the idea that Internet protocols should be applied even to the smallest devices [30].

In fact, this tiered networking architecture is not even idiosyncratic to WSN , it has been in use for long in pervasive identification systems, such as RFID, which uses a powerful *reader* to send signals to powerless *tags* within communication range to read their response. Hence, the use of this tiered architecture extends the utility of LAMP beyond the spectrum of typical sensornet applications, for example, to improve the computational capacity and transmission range of RFID and Computational-RFID (CRFID) systems. Our discussion in the remainder of this paper focuses on WSNs but the ideas presented apply just as well to pervasive identification systems.

[2]We do not use the programming model of TENET but only borrow the architectural separation of node duties in WSN

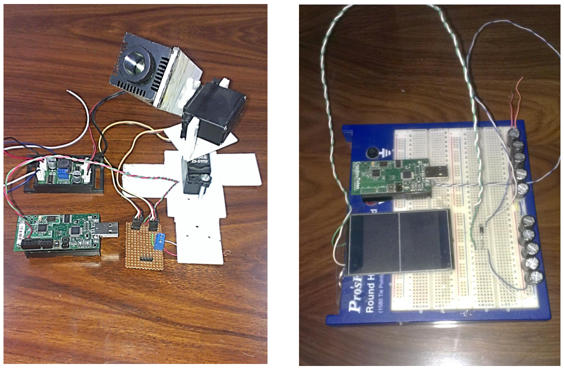

(a) LAMP master: TelosB with μ-power beaming and pan-tilt (b) LAMP leaf: TelosB with monocrystalline solar panel

Fig. 6. LAMP System Components.

We now present the detailed design of different components of our architecture.

B. LAMP Components

In the following we describe the detailed design and role of LAMP master and leaves.

1) LAMP master: A LAMP master is responsible for three functions: energy transmission, scheduling and directing energy at the appropriate leaf, and exchanging data with leaves.

Figure 6(a) shows our implementation of a LAMP master. We use laser μ-power beaming mechanism (Section III-B) to provide the energy transmission capability. We develop a dual-axis, pan-tilt mechanism using servo motors with 180° rotation possible in both axes. The laser module mounted on this platform is controlled through a TelosB mote. LAMP protocols can use this apparatus to localize and focus the laser beam on the desired leaf. Finally, the master uses the TelosB radio for data communication with leaves.

2) LAMP leaf: A LAMP leaf is required to possess two basic capabilities: the ability to receive energy, and the ability to communicate with the master and other leaves.

Our LAMP leaf is a TelosB (Figure 6(b)) equipped with a monocrystalline solar panel (same as in Section III-B) which can convert energy of a focused laser beam. We also attach a capacitor between the solar panel and input to TelosB to allow energy buffering, thus enabling duty cycled operation. If required, rechargeable batteries can also be attached but is beyond the scope of our prototype system. Using monocrystalline solar panels has an additional advantage as leaves can scavenge ambient light energy.

C. Detailed Operation of LAMP

LAMP aims to achieve three main goals: (i) provide energy transport between the energy sources (masters) and S-plane consumers (leaves), (ii) enable exchange of E-plane control information to manage energy distribution, and (iii) facilitate S-plane operation within timing and network size constraints. The overall operation of LAMP can be divided in three parts that meet these goals: bootup and searching, LAMP E-plane control protocol, and energy scheduling.

1) Bootup and Searching: Since the most important aspect of LAMP is transporting energy to nodes in-situ, a first question is for a LAMP master to know the pan-tilt coordinates for providing energy to individual leaves. We consider two possibilities for this purpose: dynamic localization with searching, or static locations known at boot-time.

Dynamic localization is possible after a fine grained search of the region assigned to each master. In this search, the master continuously shines the laser, in low power mode for safety, over its assigned region. A leaf with a power measurement circuit determines the power received; whenever this value crosses the ambient power threshold, it is reported to the master along with its ID. The master uses this positive feedback mechanism to determine and fine-tune the coordinates of a particular leaf. Thus the master will have complete location information after the search ends.

The second approach considers a static deployment, where the coordinates are manually provided to the master. The master then uses these coordinates to deliver energy to leaf nodes based on application needs and its scheduling strategy.

While we already implement support for searching in our prototype, allowing dynamic localization, we believe that optimizing the searching algorithm is a research problem on its own and part of future work. Hence, in this paper, we use the second approach to focus on evaluating the practical scalability of LAMP.

2) Control Protocol for Energy Distribution: Once a master node knows the location of a leaf, it then has to not only energize it but also exchange control information to manage energy distribution in the E-plane. We next present a first E-plane protocol that enables both these functions (Figure 7). This protocol forms the basis for a complete E-plane abstraction.

First, the LAMP master simply focuses the laser beam on a leaf node to accumulate energy on its capacitor. This charging continues until the *Response Interval* (I_{RS}) at which time the leaf has enough energy to transmit a packet. Next we describe how our protocol exchanges E-plane control information.

This protocol is a repetitive cycle defined by two types of leaf intervals: (i) *Active interval* (I_{active}) to communicate (transmit or receive) with the master, and (ii) *sleep interval* (I_{sleep}) to accumulate energy. While both these intervals are fixed in a given deployment, we later vary them to evaluate their impact on LAMP scalability.

The leaf, in the first active interval (I_{active}), transmits E-plane control information, such as requesting more energy by specifying the number of I_{sleep} intervals in which energy can accumulate. This value is at minimum two sleep intervals; during this interval the leaf turns off its radio and accumulates energy to listen for the packet to be sent by the master. This packet contains control information like the accepted number of I_{sleep} intervals, depending on the scheduling policy (see Section IV-C3). It also contains the interval for next charging epoch allowing both synchronization of each epoch between master and leaf, as well as dynamical update to the charging schedule. This charging epoch can *optionally* also pull sensed data; in fact this dual function is imperative in the clustered mode we discuss in Section IV-D. We define this variable

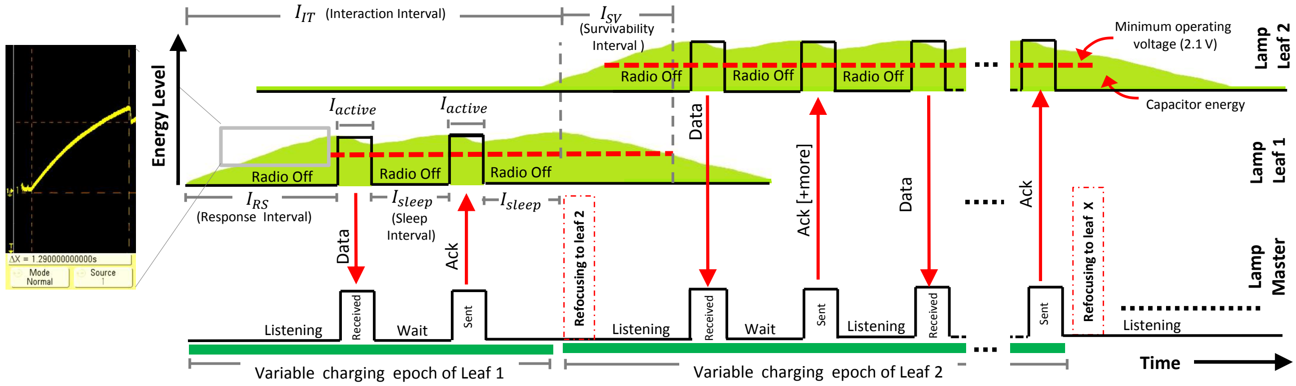

Fig. 7. LAMP Control Protocol: A master focuses laser on a leaf for a charging epoch in which leaf duty cycles to exchange control information; master then refocuses based on a scheduling policy.

charging epoch between master-leaf as the *Interaction Interval* (I_{IT}). Moreover, this repetition of I_{active} and I_{sleep} intervals constitutes the *leaf duty cycle* during any charging epoch.

Note that a leaf cannot monopolize a master since it can accept or reject the request for further I_{sleep} cycles. As described above, in case of rejection, this interaction defaults to two I_{sleep} and I_{active} cycles (leaf 1 in Figure 7).

Discussion: The tiered architecture employed by LAMP enables physical separation of energy sources and distribution (at master) from the sensing requirements of an application (at leaf). Similarly, the desired interaction between E- and S-planes is enabled by wireless energy and a variable charging epoch by allowing leaves to pull more energy. Although our prototype implementation is based on a fixed length I_{sleep} and a simple packet-based request-reply primitive for energy, it can easily support sophisticated future APIs, such as `pullenergy(...)` (Section VIII), to allow leaves to negotiate energy parameters (e.g., number and length of I_{sleep}) and resolve unequal energy consumption in a WSN. Such API can also let a leaf indicate the criticality of its demand, which can then impact the scheduling algorithm run at the master. The support for such interactions is not possible in current WSNs, with strongly coupled E- and S-planes, that treat energy as a local constraint rather than a deployment wide shareable resource as in LAMP.

3) Scheduling Energy to Meet Sensing Plane Requirements: The master, after completing interaction with one leaf, refocuses its laser beam on another leaf determined by some scheduling strategy. However, after the charging epoch, the first leaf's capacitor retains energy for some I_{SV} (survivability interval), even after the master relocates its laser beam. We use this interval to facilitate the operation of a sensing application in the S-plane; thus a traditional WSN operates as normal, with frequent charging epochs that energize it to meet its sensing requirements.

The energy scheduling policy is important in meeting these requirements. Similar to managing the demand and supply on the power grid, our energy scheduler is meant to assign charging epochs to leaves in order to best meet the application requirements. The scheduler has to do so within the system constraints of not only supply but in our case, also timing, as power is beamed individually at some rate. In our current evaluation we schedule energy network wide in a round-robin fashion where we cycle through a spatially ordered list of leaves. with master periodically visiting each leaf at the LAMP recharging rate (LRR). For this scheduling our evaluation shows that to maintain a fully active WSN of a particular size, our system mechanics introduce strict timing requirements on LRR (Section V-D).

We can see more involved scheduling policies that depend on a wide range of aspects: from static priorities like Earliest *Dead* First, to application-driven schemes that meet timing constraints. However, we defer the choice of scheduling policies to future work.

D. Supported Communication Modes

We now highlight the two communications modes enabled in the S-plane by our E-plane LAMP architecture. Both these communication modes offer different advantages while imposing restrictions on performance characteristics such as LRR and the number of leaves per master (i.e., scalability).

Cluster Mode: Leaves only interact with each other through their respective masters. Here we use the packets in an I_{INT} to purposefully exchange sensed data (as opposed to the default for only E-plane control information). This mode is similar to the infrastructure mode in IEEE 802.11 where the hosts can only interact through their access point. The principal advantage of this communication mode is that it improves scalability of LAMP (cf. Section V-D). This improvement is because a leaf can live longer by turning its radio off once the interaction with master is over, allowing the master to serve more leaves before returning.

Mesh Mode: Leaves are free to directly communicate with each other by using the overlap among I_{SV} of leaves. This mode is similar to the ad hoc mode in IEEE 802.11. The principle advantage of this mode is that it enables LAMP to support typical WSN applications requiring a flat, un-tiered node deployment. It is clear that this mode is more power hungry and less scalable.

However, in both modes scalability is only limited if we want an *active network* where all the leaves are recharged before their power levels drop below the operational threshold of TelosB. Otherwise, the scalability of this mode is only affected by application constraints such as the desired recharging- and data-rates. We provide a detailed qualitative analysis for these different communication patterns in Section V-D.

V. LAMP Evaluation

Having explained the implementation and operation of our LAMP system, we now investigate its practical limits. We thus answer two fundamental questions: what is the maximum application data rate and duty-cycle supported by LAMP? What is the scalability of LAMP in terms of the number of leaves per master? The answers to these questions allow us to fully understand the overall benefits as well as limitations of our system.

To answer these questions we first need to understand the micro-characteristics of the LAMP master and leaf hardware.

A. Master Movement Characterization

The characteristics of the servo-based pan-tilt mechanism at the master, and the charging and discharging dynamics of the leaf, constrain the overall LAMP system. We first characterize these micro-aspects.

1) Minimum node separation: The first micro-aspect of our system we establish is the minimum separation between nodes to allow for proper laser focusing. For this purpose we evaluate the minimum resolution, in degrees, of the servo motors used to move the mounted laser. While information about resolution and speed is provided in a data sheet, it is for an unloaded motor. We control the mechanism using two Futaba s3003 servo motors over the ADC1 and ADC2 channels on the TelosB. For measuring the resolution we vary the PWM signal to these motors in increments of 8μs, and observe the laser point on a wall 15m away. We measure the first change in the location of the laser point, and then, using trigonometry, compute angular change along a single axis. We repeat this experiment three times and compute the average which turns out to be 0.49° along any axis.

This minimum resolution gives us a minimum separation of 0.87m for leaves at 100m of radial distance from the master, an acceptable inter-node separation.

2) Limits on refocusing time: We now find the minimum and maximum time it takes for a master to refocus from one leaf to another. We use this result in computing the bounds on the LRR. To measure these limits, we simply measure the power drawn by the motor as we command it to move the full 180° arc. Figure 8(b) shows the measured transition taking 680ms (max refocusing time), implying that it takes 3.78ms for a single degree of movement. To find the minimum refocusing time we observe that our smallest possible movement of 0.49° requires just 1.85ms.

3) Calibration requirements: For recharging the leaf, we need to ensure that the laser-beam correctly points at the solar panel of the leaf. The master might require to re-calibrate once it returns to the same leaf after recharging all other assigned leaves. The frequency and the magnitude of this re-calibration strongly depends upon the movement-error characteristics of a particular servo motor.

We measure this error for our servo motors by fixing the laser beam at a certain point, i.e., the origin, before moving it a complete 180° along one axis and then back. After each iteration we manually measure (with a metric rule) the radial distance from the origin. We repeat this experiment 100 times. Figure 8(a) shows the CDF of these errors from the origin. We can clearly observe that the errors are *bounded* with a maximum radial error of 4.5cm, preventing cumulative error build up. Also, nearly 90% of all error values lie between 3-4.5cm. During our experiments, we also observe that these errors do not self-cancel as they remain in the same quadrant.

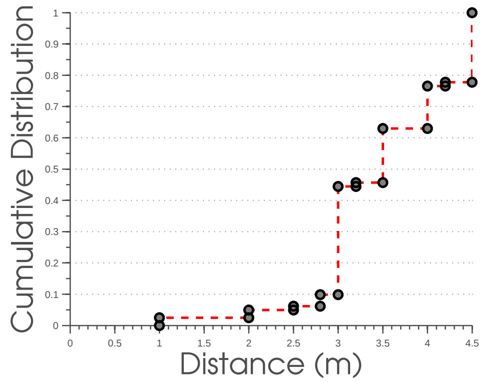

(a) CDF of servo motor error

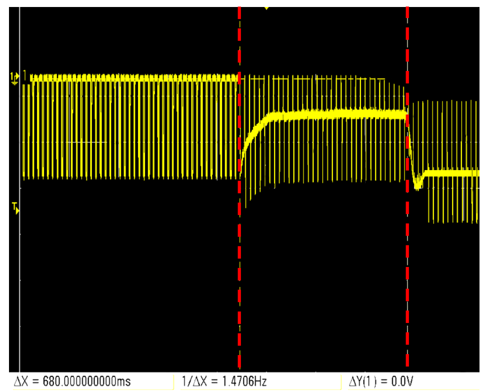

(b) Servo motor speed

Fig. 8. Servo motor characteristics.

Thus, for a leaf node with 8.8×5.4 cm solar panel and a perfectly centered laser beam, in the worst case, the master will have to recalibrate after every round trip if the master-leaf distance is greater than 29.5m.

Technological Limitations: We emphasize that these error bounds are for a particular motor technology, i.e., Futaba's s3003. Hence, these results should not be generalized to define the accuracy limits of laser-beaming used in an energy distribution architecture including LAMP. The selection of a motor technology is a tradeoff between its cost and the required re-calibration effort for a particular architecture. Since re-calibrating the master is not within the scope of this prototypical study, our choice is biased towards the overall cost concerns of LAMP.

B. Charging Behavior of the LAMP leaf

The choice of capacitor determines the charging and discharging dynamics of a LAMP leaf. A Capacitor is like a water reservoir; the voltage across a capacitor is determined by the

TABLE II. LEAF I_{RS} WITH DIFFERENT CAPACITORS.

Capacitor (μ**F**)	100	2000	4000	6000	8000	1000
I_{RS} **(seconds)**	1.29	6.58	10.91	15.74	23.15	26.93

TABLE III. I_{IT} FOR TWO PACKET INTERVALS AS WE VARY CAPACITANCE AND SLEEP INTERVAL.

	I_{sleep} **(seconds)**					
Capacitor (uF)	**0.1**	**1**	**2**	**3**	**4**	**5**
100	1.498	3.298	5.298	7.298	9.298	11.298
2000	6.786	8.586	10.586	12.586	14.586	16.586
4000	11.114	12.914	14.914	16.914	18.914	20.914
6000	15.943	17.743	19.743	21.743	23.743	25.743
8000	23.36	25.16	27.16	29.16	31.16	33.16
10000	27.137	28.937	30.937	32.937	34.937	36.937

level of charge. For a fixed rate of charging, this charge level is reached later for a larger capacity of reservoir (its capacitance). Thus, while a smaller capacitance quickly reaches a defined usable voltage, it delivers lesser amount of current.

With this basic understanding we experimentally evaluate two important system parameters; the response interval I_{RS} in which a leaf responds after the master focuses the laser beam on it, and the total interaction interval I_{IT} (see Figure 7).

1) How quickly can a leaf be energized?: We measure the response interval by focusing the laser beam on a completely discharged leaf and use an oscilloscope to observe voltage across a capacitor bank that we use to vary capacitance. The sudden changes in the voltage level on the oscilloscope clearly indicate the charging of the capacitor and the subsequent transmission of the first packet (see the scope capture in Figure 7), allowing us to measure I_{RS}.

Table II shows the I_{RS} for different capacitance values. An important observation is that our TelosB-based leaf requires at least 100μF capacitance to run the LAMP protocol. Lower valued capacitors simply cannot store enough energy to allow the initial transmission. As we increase the capacitor size, a correspondingly linear increase in I_{RS} is observed. Thus to minimize I_{RS} we should simply use this minimum capacitance.

2) What is the minimum interaction interval? : We now want to understand the dynamics of the master-leaf interaction interval (I_{IT}) . Using Figure 7, the I_{IT} of a leaf can be derived as follows:

$$I_{IT} = I_{RS} + (I_{active} + I_{sleep}) \times K \quad (1)$$

Where K is the number of equal duration (I_{active} and I_{sleep}) intervals lower bounded to two in the LAMP protocol (described in Section IV-C2). The sleep interval I_{sleep} is for charging a leaf; this interval is fixed for a given deployment, but we vary it here to understand its impact on scalability. We fix I_{active} to 8ms based on results reported for the CC2420 radio on TelosB [3]. Transmission failure during this interval, for example, due to MAC contention, can be handled by increasing the I_{IT} as shown in Figure 7 for leaf 2.

With the minimum I_{RS} defined in the previous section, the shortest I_{IT} is therefore now a function of the I_{sleep}. We measure this interval using a small experiment: we set a leaf to transmit an increasing counter value in I_{active}. We keep the laser beam focused on the leaf while decreasing I_{sleep}; we infer its smallest value when we start observing periodic reset of the counter, indicating that the leaf lost power before its listen interval.

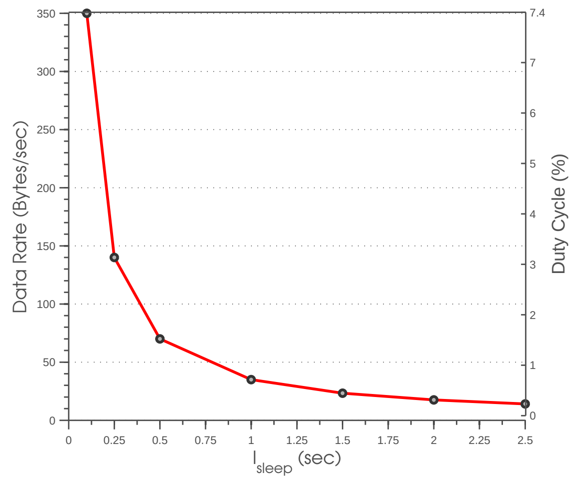

Fig. 9. LAMP Application throughput: varying I_{sleep} and its duty cycle.

We find that $I_{sleep} = 100ms$ is the smallest interval, regardless of the capacitor size. This observation is explained by noting that the transmit and receive energy cost, for same packet size, are comparable for CC2420 [28]. Thus the leaf must only recover the same amount of energy lost during its transmission to receive the subsequent packet from the master. Thus, increasing capacitor beyond the the minimum required for transmitting that first packet does not affect the I_{IT}, provided that enough charge is recovered during I_{sleep}.

Table III shows the interaction interval for different values of I_{sleep} and different capacitor sizes, simply by plugging in empirical values of I_{RS} (which is a function of capacitor size) into Equation 1.

Discussion: Our evaluation of the energy dynamics of the leaf indicate that for the fastest interaction we should use the smallest possible capacitor size of 100μF with the smallest I_{sleep} of 100msec. However, we also note that, larger capacitor and I_{sleep} impact the I_{sv} that then determines the scalability of LAMP in terms of supported leaves/master and LRR. We evaluate these factors in Section V-D.

With these system characteristics defined, we next answer our two questions that will help define the practical limits of our LAMP implementation.

C. Application Duty-cycle and Data Rate

A first important question we want to answer is about the data rate during master-leaf interaction and the largest leaf duty cycle supported by LAMP. The duty cycle in this particular scenario is the duration for which the radio hardware of leaf remains active during I_{IT}.

An answer to both these question helps understand the range of WSN applications that can be supported by our system. Using laser μ-power beaming mechanism, we expect to support a duty cycle of around 10% (cf. Figure 4).

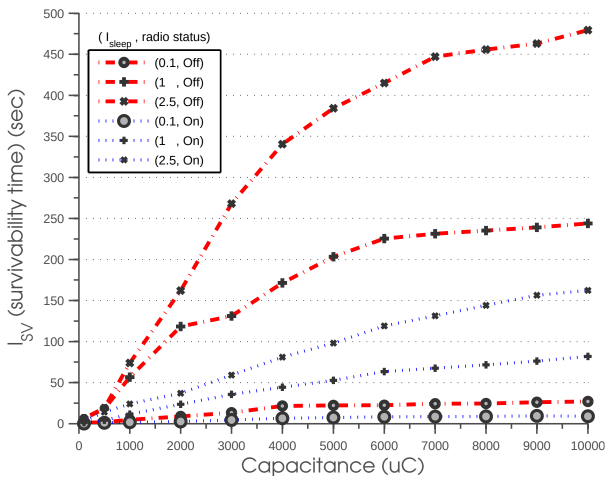

Fig. 10. I_{SV} as a function of I_{sleep} (for recharging), and capacitance (charge capacity).

We vary I_{sleep} to alter the leaf's duty cycle. Figure 9 shows that LAMP can provide a maximum data rate of 350Bps at a duty cycle of 7.4%, coming quite close to the 10% we expected after our power harvesting experiments in Section III. This small difference is attributed to harvesting inefficiencies of TelosB since its resistance is not optimized for the solar panel.

D. LAMP Scalability

We finally evaluate the scalability of LAMP in terms of the number of leaves per master it supports, and any tradeoff with the required LRR. The scalability of LAMP is strongly dependent upon I_{SV} which is a function of charge retained by the capacitor even after the master has refocused to another leaf. Thus, we first empirically observe the impact of different factors on I_{SV}. We then describe the model employed to determine the scalability of our system. We finally use our measurements and model to thoroughly understand the scalability limits.

1) Survivability Limits for leaves: I_{SV} is a function of three parameters: capacitor size, I_{sleep}, and whether we periodically use radio (e.g., in mesh mode) that dominates the energy consumption of TelosB. We measure this time empirically by observing the capacitor voltage dropping below 2.1V (the minimum operational voltage for TelosB), while we vary these parameters. Figure 10 depicts the result and we can clearly see that increasing I_{sleep} significantly improves I_{SV} when using larger capacitors, in both on and off radio states. Thus, only when a capacitor is large enough to store the excess energy collected over a longer sleep interval do we see greater survivability.

2) Modeling Recharging Rate: We observe that a given survivability time can be utilized to either increase the number of leaves (N) that can be served or alternatively increase the LRR (how often master returns to the same leaf). In order to understand the tradeoff between these two variables, we model a network where leaves are kept equidistant. The following equation then defines LRR based on the interaction time (I_{IT}), number of leaves (N), and $T(X)$ which is movement time required by the servo motor to move X degrees between leaves.

$$LRR = N \times I_{IT} + 2(N-1) \times T(X) \tag{2}$$

If we want to maintain an always *active network*, the master must return to recharge each leaf before its power level falls below 2.1 V, which, by definition, happen after I_{SV}. Note that for such a scenario, $I_{RS} = 0$, significantly reducing the interaction interval with each leaf. In order to maximize the number of leaves supported in any setting, we make two observations. First, since a larger LRR supports greater number of leaves, we use the largest possible value of $LRR = I_{SV}$. Similarly, we also keep all leaves separated by 0.49°; the minimum separation allowed by our servo motor (Section V-A). Plugging the value of $LRR = I_{SV}$ from our experiments (Figure 10) and the constant value for $T(X) = 1.85ms$ allows us to determine the recharging rate for a particular number of leaves per master for different configurations of I_{IT}.

3) Results: In an active network, the communication mode (Section IV-D) greatly impacts the duration of I_{SV}. We therefore separately evaluate both the mesh and clustered mode of operation.

Recall that for **mesh operation**, all leaves use their radio for direct communication with each other even after the master relocates its laser beam. While any radio duty cycle can be supported, for simplicity we continue with the duty cycle determined by I_{sleep} for our evaluation. We notice that LAMP can scale up to 40 leaves per master with a recharging rate of 9.50 seconds in this mode (Figure 11(a)). This implies that for a LAMP recharging rate of 9.5s, with a 10mF capacitor, we can *perpetually* energize a 40 node, full-mesh, WSN.

We next look at the **cluster operation**, where leaves only *sense* at an application defined rate with no intra-leaf communication after the master relocates its laser beam (all data routing occurring through the master during a LAMP interaction). With lesser energy consumption we expect it to achieve greater scalability than mesh operation. Figure 11(b) validates this as LAMP can now scale up to 120 leaves per master with a recharging rate (or polling rate) of 26.5s. This implies that for a LAMP recharging rate of 26.5s, again with a 10mF capacitor, we can *perpetually* energize a 120 node WSN that networks only via masters.

One observation across both mesh and cluster modes is that LAMP scales better for smaller I_{sleep}. This correlation is for two reasons: First, a capacitor's non-linear charging curve results in smaller charge accumulation after its time constant, thus most of the charge for survivability is achieved quite early in a sleep interval. Second, having a larger I_{sleep} increases the interaction interval, further negatively impacting the number of leaves supported. We also observe minor variations in scalability seen for Figure 11(a) and Figure 11(b) due to the very dynamic discharging behavior of capacitors [11] which we do not model but is captured in the measurement of the survivability time.

Finally, if we want to support greater number of leaves we need to increase the LRR to greater than its survivability time. Thus, leaves will shutoff, and only disruption tolerant WSN applications can benefit from this scenario. Such applications can then support an arbitrary number of leaves but with a corresponding increase in the LRR, which determines the

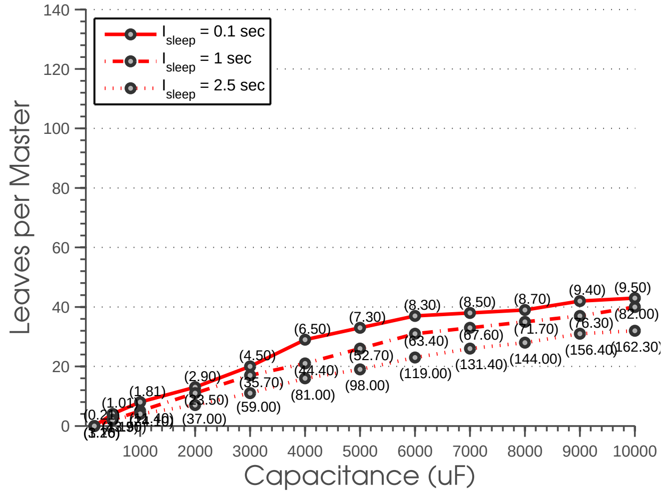

(a) Mesh: Radio use between leaves

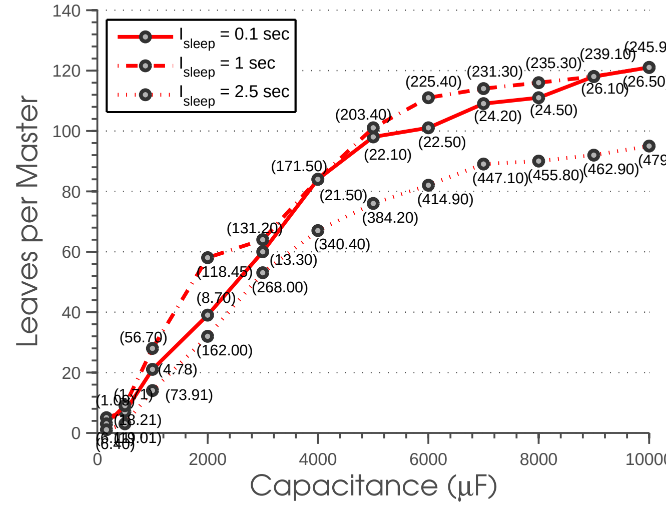

(b) Cluster: Radio use only during master-leaf interaction

Fig. 11. While maintaining an always-on network, LAMP scales up to 40 and 120 leaves per master in mesh and cluster modes, respectively (marker values represent recharging rate in seconds).

duration of disruption. These results can be generated directly from Equation 2; however we now include the response time which introduces a significant increase in recharging rate. Thus, with response time included, supporting 200 nodes in cluster mode needs LRR of 300.5s, in the best case, instead of 28.5s we needed to support 120 leaves.

Overall, we can conclude that scalability of LAMP is well suited for practical WSN deployments, and can be controlled by user-controlled parameters such as recharging rate, capacitor size and sleep interval.

VI. Limitations

Our work is the first attempt to decouple energy from sensing through a laser-based power beaming solution. As such, we realize it has some practical limitations that we identify with their possible solutions, as future work, below.

Line-of-sight (LoS): LAMP architecture currently requires LoS between a master and a leaf. We believe that adding a third node-type, having mirrors, in our architecture can allow a power-beam to route around obstacles. We see this achievable with previous work for outdoor settings on reflecting sunlight [18], [26], and indoors using mirrors on the ceiling as used for a non-LoS 60Ghz wireless communication in datacenter networks [36].

Localization and (Re)calibration: The practical scalability of our proposed system depends largely on the calibration and localization accuracy. While we do meticulously identify the calibration limits for our pan-tilt mechanism (Section V-A), we agree that this issue needs further investigation for a fully functional deployment. We believe that technology for such precise aiming exists, with laser power-beaming used on mobile UAVs 1km away with cm-level accuracy [29], and the high cost precluding its current use in our system, will reduce as the technology matures. We thus believe efficient localization of leaf nodes is a fertile area of future research.

Safety: A final area of concern are the safety issues when using lasers. We show that by using just a 1W laser, its is feasible to support low-power mote-class devices, although eye-safety and long term exposure still are areas of concern. We can increase the safety by using "eye-safe" lasers at wavelength beyond 1500nm [15]. Furthermore, we can use a very low-power (few mW) pilot beam while localizing the leaf node and use the RF channel for feedback to move to full power transfer once aligned. This same channel can be explored to provide feedback such that when the powerbeam is broken (e.g., by animals or humans) the energy beam can immediately be switched off.

Despite these limitations, this paper is a significant step forward in identifying the limits and scales of a practical wireless-energy transfer mechanism for WSNs.

VII. Related Work

The scope of our paper is quite novel; to our knowledge, LAMP is the first system that uses laser μ-power beaming to build an energy distribution architecture for WSN. Unlike theoretical formulations [2], [6], this paper emphasizes the practical feasibility of a wireless energy transfer mechanism and thoroughly evaluates its limits to establish its feasibility boundaries. Moreover, unlike earlier work using light [18], [26], LAMP is not just an energy transfer mechanism; it is a complete architecture providing the abstraction of an energy plane that can transparently integrate with *any* existing sensing architecture to provide energy as a service.

We can broadly divided the related research in three main categories.

Physical Transfer: The naive solution for energy transfer is to physically carry energy through the network [17] and recharge node batteries. The research in this area mainly focuses on an optimization problem of finding the best path through the network that results in recharging maximum nodes in minimum time. However, physical transfer of energy requires human-driven or robotic vehicular movement across the network strongly limiting deployment scenarios.

Ambient Harvesting: A great deal of research deals with ambient harvesting of energy to buy longer life times both in WSN [13], [21], [35] and in pervasive identification systems [11]. However, energy harvesting requires favorable conditions and thus restricts spatial deployment of nodes only at energy rich areas. Similarly, energy harvesting alone is not sufficient to achieve our main goal, i.e., to decouple energy and sensing planes in WSN. For this purpose our primary focus in this paper is to utilize abundant energy sources nearby a deployment to develop an efficient wireless energy transfer mechanism that can nonetheless support additional sources of ambient energy.

Light Reflection: Approaches such as elighthouse [18] and others [26] propose reflecting sun light on to sensor nodes using mirrors. However, as we demonstrate in Section III-B3, the range and magnitude of such energy transfer is insufficient for running a mote-class device such as TelosB indefinitely even with LoS between the light source and the node.

VIII. Conclusions and Future Directions

In this paper we propose to decouple energy and sensing plans in WSN to utilize abundant energy sources distant from the sensing location allowing unrestricted network lifetime. We compare different wireless energy transfer mechanisms to identify that laser μ-power beaming satisfies the fundamental goals of cheap, in-situ, collaborative, and long lived sensing in WSN. We also present the first energy distribution architecture that enables decoupling energy and sensing planes. Our results demonstrate the feasibility of such an architecture in terms of data rate, duty-cycle, and scalability.

The argument for an energy plane abstraction opens up a great body of intriguing future work. Much like the early era of WSN, we envision research on both hardware and software aspects of small scale energy distribution architecture. On the software side, we expect refinement in the basic API that sensing applications (in the S-plane) will use to request energy (in the E-plane) through an energy transfer technology. Similarly, building the entire network stack for the E-plane elements, including an optimized version of our current LAMP protocol, is a promising future direction. On the hardware side, we see mobile and pervasive computing demands push the envelope of wireless energy technology. We see this being used to build more compact and low-cost technologies, with sensors being able to measure their energy "RSSI". Furthermore, employing different technology for energy distribution will require a different implementation of the energy plane abstraction. We even see potential in modulating the energy transfer to enable communication between devices, thus precluding the need of using the expensive radio during the interaction interval.

Acknowledgment

Many thanks to our Shepherd, Bodhi Priyantha, for his help in improving the quality of this paper. We would like to thank and acknowledge Ali Khayam and Adnan Iqbal, who were involved in earlier insightful discussions on decoupling energy and sensing planes in WSN. We also thank Osama Haq and Omer Naeem for their help with early prototypes and experimentation. This work was partially (in equipment) funded by the NGRI initiative of the Pakistan National ICT R&D fund, and the FAST-NU research fund.

References

[1] 3W CREE Q3 LED flashlight. http://tinyurl.com/3WFlashlight.

[2] M. I. Afzal, S. Seoyong, W. Mahmood, and S. M. Sajid. Optical wireless recharging model for wireless sensor nodes by using corner cube retroreflectors (ccrs). International Conference on Hybrid Information Technology (ICHIT) '09, pages 626–631, New York, NY, USA, 2009. ACM.

[3] M. H. Alizai, O. Landsiedel, J. A. B. Link, S. Götz, and K. Wehrle. Bursty traffic over bursty links. In *Proceedings of the 7th ACM Conference on Embedded Networked Sensor Systems*, SenSys '09, pages 71–84, New York, NY, USA, 2009. ACM.

[4] G. Barrenetxea, F. Ingelrest, G. Schaefer, and M. Vetterli. The hitchhiker's guide to successful wireless sensor network deployments. In *Proceedings of the 6th ACM SenSys Conference*, pages 43–56, Raleigh, NC, USA, Nov. 2008. ACM.

[5] T. Blackwell. Recent demonstrations of Laser power beaming at DFRC and MSFC. In A. V. Pakhomov and L. N. Myrabo, editors, *Beamed Energy Propulsion*, volume 766 of *American Institute of Physics Conference Series*, pages 73–85, Apr. 2005.

[6] R. Doost, K. Chowdhury, and M. Di Felice. Routing and link layer protocol design for sensor networks with wireless energy transfer. In *Global Telecommunications Conference (GLOBECOM 2010), 2010 IEEE*, pages 1–5, 2010.

[7] D. Estrin, R. Govindan, J. Heidemann, and S. Kumar. Next century challenges: Scalable coordination in sensor networks. In *Proceedings of the 5th Annual ACM/IEEE International Conference on Mobile Computing and Networking*, MobiCom '99, pages 263–270, New York, NY, USA, 1999. ACM.

[8] D. Ganesan, B. Greenstein, D. Perelyubskiy, D. Estrin, and J. Heidemann. An evaluation of multi-resolution storage for sensor networks. In *Proceedings of the 1st International Conference on Embedded Networked Sensor Systems*, SenSys '03, pages 89–102, New York, NY, USA, 2003. ACM.

[9] O. Gnawali, B. Greenstein, K.-Y. Jang, A. Joki, J. Paek, M. Vieira, D. Estrin, R. Govindan, and E. Kohler. The Tenet architecture for tiered sensor networks. In *Proceedings of the Fourth ACM SenSys Conference*, pages 153–165, Boulder, Colorado, USA, Nov. 2006. ACM.

[10] M. Gorlatova, P. Kinget, I. Kymissis, D. Rubenstein, X. Wang, and G. Zussman. Challenge: ultra-low-power energy-harvesting active networked tags (EnHANTs). In *In Proceedings of the 15th annual international conference on Mobile computing and networking Mobicom '09*, Sept. 2009.

[11] J. Gummeson, S. S. Clark, K. Fu, and D. Ganesan. On the limits of effective hybrid micro-energy harvesting on mobile crfid sensors. In *Proceedings of the 8th International Conference on Mobile Systems, Applications, and Services*, MobiSys '10, pages 195–208, New York, NY, USA, 2010. ACM.

[12] IXYS. SolarMD SLMD481H08L. web site http://ixapps.ixys.com, 2007.

[13] A. Kansal, J. Hsu, S. Zahedi, and M. B. Srivastava. Power management in energy harvesting sensor networks. *ACM Trans. Embed. Comput. Syst.*, 6(4), Sept. 2007.

[14] A. Kurs, A. Karalis, R. Moffatt, J. D. Joannopoulos, P. Fisher, and M. Soljacic. Wireless power transfer via strongly coupled magnetic resonances. *Science*, 317(5834):83–86, July 2007.

[15] Laser Focus World. EYE-SAFE LASERS - Retina-safe wavelengths benefit open-air applications. http://tinyurl.com/laserfocusworld.

[16] P. Levis, D. Gay, and D. Culler. Active sensor networks. In *Proceedings of the 2Nd Conference on Symposium on Networked Systems Design & Implementation - Volume 2*, NSDI'05, pages 343–356, Berkeley, CA, USA, 2005. USENIX Association.

[17] K. Li, H. Luan, and C.-C. Shen. Qi-ferry: Energy-constrained wireless charging in wireless sensor networks. In *Wireless Communications and Networking Conference (WCNC), 2012 IEEE*, pages 2515–2520, 2012.

[18] P. Liu, Y. Wu, J. Qiu, G. Dai, and T. Fu. elighthouse: Enhance solar power coverage in renewable sensor networks. *International Journal of Distributed Sensor Networks*, 2013.

[19] S. R. Madden, M. J. Franklin, J. M. Hellerstein, and W. Hong. Tinydb:

an acquisitional query processing system for sensor networks. *ACM Trans. Database Syst.*, 30(1):122–173, Mar. 2005.

[20] J. Paek, B. Greenstein, O. Gnawali, K.-Y. Jang, A. Joki, M. Vieira, J. Hicks, D. Estrin, R. Govindan, and E. Kohler. The tenet architecture for tiered sensor networks. *ACM Transactions on Sensor Networks (TOSN)*, 6(4):34:1–34:44, July 2010.

[21] C. Park and P. Chou. AmbiMax: Autonomous energy harvesting platform for multi-supply wireless sensor nodes. volume 1, pages 168–177, Sept. 2006.

[22] PowerCast. P2110 receiver and TX91501-3W-ID transmitter. http://tinyurl.com/powercaster, 2009.

[23] J. Rodrigues, O. Pereira, and P. Neves. Biofeedback data visualization for body sensor networks. *J. Netw. Comput. Appl.*, 34(1):151–158, Jan. 2011.

[24] T. Schmid, R. Shea, M. B. Srivastava, and P. Dutta. Disentangling wireless sensing from mesh networking. In *Proceedings of the 6th Workshop on Hot Topics in Embedded Networked Sensors*, HotEmNets '10, pages 3:1–3:5, New York, NY, USA, 2010. ACM.

[25] F. Silva, E. A. Graham, A. DeSchon, J.-H. Chang, Y. Cho, P. W. Rundel, and T. Benzel. Span: A sensor processing and acquisition network - field deployment lessons learned. In *94th ESA Annual Meeting*, 2009.

[26] A. A. Syed, Y. Cho, and J. Heidemann. Energy transference for sensornets. In *Proceedings of the 8th ACM Conference on Embedded Networked Sensor Systems*, SenSys '10, pages 397–398, New York, NY, USA, 2010. ACM.

[27] J. Taneja, J. Jeong, and D. Culler. Design, modeling, and capacity planning for micro-solar power sensor networks. In *Proceedings of the 7th International Conference on Information Processing in Sensor Networks*, IPSN '08, pages 407–418, Washington, DC, USA, 2008. IEEE Computer Society.

[28] Texas Instruments. Chipcon smartRF CC2420. http://tinyurl.com/CC2420datasheet.

[29] Thomas Nugent,Jordin Kare,David Bashford,Carsten Erickson,Jeff Alexander. 12-hour hover: Flight demonstration of a laserpowered quadrocopter. White Paper, 2011.

[30] J.-P. Vasseur and A. Dunkels. *Interconnecting Smart Objects with IP: The Next Internet*. Morgan Kaufmann Publishers Inc., San Francisco, CA, USA, 2010.

[31] C. M. Vigorito, D. Ganesan, and A. G. Barto. Adaptive control of duty cycling in energy-harvesting wireless sensor networks. In *SECON*, pages 21–30. IEEE, 2007.

[32] Wuhan Lilly Electronics Co. Focusable 0.8W near-infrared IR laser diode module. web site http://tinyurl.com/0-8WLaser.

[33] L. Xie, Y. Shi, Y. T. Hou, and H. D. Sherali. Making sensor networks immortal: An energy-renewal approach with wireless power transfer. *Networking, IEEE/ACM Transactions on*, 20(6):1748–1761, 2012.

[34] W. Ye, F. Silva, and J. Heidemann. Ultra-low duty cycle MAC with scheduled channel polling. In *Proceedings of the 4th International Conference on Embedded Networked Sensor Systems*, SenSys '06, pages 321–333, Boulder, Colorado, USA. ACM.

[35] C. Zhang, A. Syed, Y. Cho, and J. Heidemann. Steam-powered sensing. In *Proceedings of the 9th ACM Conference on Embedded Networked Sensor Systems*, SenSys '11, pages 204–217, New York, NY, USA, 2011. ACM.

[36] X. Zhou, Z. Zhang, Y. Zhu, Y. Li, S. Kumar, A. Vahdat, B. Y. Zhao, and H. Zheng. Mirror mirror on the ceiling: Flexible wireless links for data centers. In *Proceedings of the ACM SIGCOMM 2012 Conference on Applications, Technologies, Architectures, and Protocols for Computer Communication*, SIGCOMM '12, pages 443–454, New York, NY, USA, 2012. ACM.

[37] T. Zhu, Y. Gu, T. He, and Z.-L. Zhang. eShare: A Capacitor-Driven Energy Storage and Sharing Network for Long-Term Operation. In *sensys'10*, pages 239–252.

[38] T. Zhu, Y. Gu, T. He, and Z.-L. Zhang. Achieving long-term operation with a capacitor-driven energy storage and sharing network. *ACM Trans. Sen. Netw.*, 8(4):32:1–32:37, Sept. 2012.