An Autonomous Distributed Demand-side Energy Management Network Using Fluorescent Lamp Sensors

John J. Cooley[*], Al-Thaddeus Avestruz[**], Steven B. Leeb[***]

[*]IEEE Student Member, MIT Rm. 10-007, Cambridge, MA 02139 - jjcooley@mit.edu
[**]IEEE Member, MIT Rm. 10-017, Cambridge, MA 02139 - avestruz@mit.edu
[***]IEEE Fellow MIT Rm. 10-069, Cambridge, MA 02139 - sbleeb@mit.edu

Abstract - **This paper demonstrates the building blocks of a new autonomous demand-side energy management system. The system uses a network of lamp sensors interfaced with frequency-controlled dimming ballasts. The auto-dimming lamps detect occupants and automatically adjust their own brightness levels. These occupancy detections may be used to adjust power consumption of other systems such as heating, ventilation and air-conditioning (HVAC) as well. This paper also demonstrates a quasistatic frequency-modulated (FM) wireless link to network auto-dimming lamps with each other. The wireless link reuses the frequency-controlled action of the dimming ballast to transmit data and the lamp sensor electric field measurement to receive data. All of the electronics are made to fit inside a ballast box with the intent to create a drop-in replacement for standard ballasts. This paper presents critical system characterizations and design strategies as well as a demonstration of an auto-dimming lamp and the wireless link.**

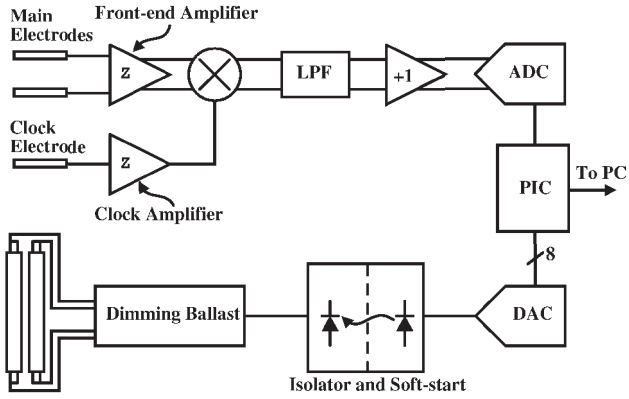

Fig. 1. The block diagram of the auto-dimming lamp sensor. Transimpedance amplifiers are marked with a 'Z'.

I. Introduction

There is a great interest in controlling lighting to optimize energy consumption. Lighting in commercial and residential spaces consumes a significant portion of the end use demand for delivered energy in the United States. In 2005, lighting consumed 0.73 Quadrillion Btu (QBtu) in the residential sector and 1.18 QBtu in the commercial sector [15]. This accounts for 15.6% and 13.9% of the total electricity delivered in the residential and commercial sectors, respectively; approximately $20.1 billion and $29.7 billion spent by electricity consumers in the residential and commercial sectors respectively [15].

The U.S. Environmental Protection Agency (EPA) encourages reduction of energy consumption by improving efficiency of energy systems such as lighting. The EPA's Energy Star program provides energy efficient solutions for reducing energy consumption while maintaining or improving the current standards of living [17]. The Energy Star program also recognizes businesses and organizations for reducing greenhouse gas emissions through energy efficiency. "In 2006 alone, Americans with the help of Energy Star saved $14 billion on their energy bills and avoided greenhouse gas emissions equivalent to those of 25 million vehicles [17]." According to the EPA's Energy Star 2006 Annual report, Energy Star lighting saved 11.3 Billion kilawatt-hours(kWh) in 2006. This eliminated 2.2 Million Metric tonnes of Carbon Equivalent (MMTCE) of pollution in the U.S. that year [18].

Many references are available on demand-side energy management by controlling loads such as lighting. Reference [12] discusses an internet-enabled energy management system motivated by large energy consumers' need for "efficient, automated ways of switching or adjusting noncritical loads (lighting, heating, air-conditioning, chilled water, etc.) [12]" The autonomous energy management system in this work can adjust lighting levels and it could also be used to monitor activity in order to adjust other "noncritical loads" such as heating and air-conditioning.

Reference [13] is one of many references that discusses curtailed demand and its value in the energy market. In curtailed demand, the supply company reduces the effective energy demand by reclaiming unused or wasted energy. Reference [13] cites the sophisticated planning and knowledge of building characteristics necessary to implement curtailed demand and suggests that it can only be effective if "sensing and switching can be done cheaply" and with "a high level of automation."

A recent article in the *New York Times* also discusses approaches to curtailed demand. Enernoc Co., a company in Boston, can remotely manage demand-side energy consumption in businesses nationwide to curb demand during high demand times [14].

In [10], the authors discuss the application of micro-electromechanical (MEMS) illuminance sensors spread

978-1-4244-1668-4/08/$25.00 ©2008 IEEE

around work spaces to measure daylight. The MEMS sensors are intended to be interfaced with dimming fluorescent lamp ballasts to adjust lighting levels based on the daylight measurements.

This paper presents the building blocks of a new autonomous and self-expanding demand-side energy management system for lighting control. It is a demand-side energy management system because the control of energy consumption is confined to the demand or end-user portion of the power system. References [1]–[4] detail the design of the "lamp sensor" which measures the lamp's own electric fields to detect human targets. Interfacing the lamp sensor with a dimming ballast creates a smart auto-dimming lamp which uses the lamp sensor's occupancy detections to appropriately dim or brighten. These fine-grain occupancy detections may also be used to adjust power consumption of other systems such as Heating, Ventilation and Air Conditioning (HVAC). All of the electronics are made to fit inside a ballast box with the intent to create a drop-in replacement for standard ballasts. We have also demonstrated a quasistatic frequency-modulated (FM) wireless link to enable communication between adjacent lamps. By communicating with adjacent lamps, one lamp can command a cluster or an entire room of lamps to turn on above an occupant according to desired lighting schemes. The wireless link reuses the lamp sensor electric field measurement and the action of the frequency-controlled dimming ballast.

Because the performance of the lamp sensor system depends on many factors such as ceiling height, electrode configuration, separation of lamps, the physical geometry of the particular lamp case, etc., the auto-dimming lamp in this paper is presented as an example. From this example, we detail critical system characterizations that must be understood by the system designer to implement a network of auto-dimming lamps.

Section II presents electronics including a dimming ballast, an updated lamp sensor and its interface to the ballast. Section III presents system characterizations necessary to implement the auto-dimming lamps including the lamp as a signal source and the lamp sensor signal-to-noise ratio (SNR) across dimming levels. It also presents example characterizations of three critical non-idealities in the lamp sensor system: drift, offsets and offset settling after switching between dimming levels. It then presents a strategy for using this information to design an auto-dimming lamp including auto-calibration and detection strategies as well as a demonstration of a real auto-dimming lamp. Section IV presents a novel wireless link between adjacent lamps including the design, simulation and measurement of a phase-locked loop (PLL).

II. Dimmed Sensing: Electronics and the Signal Source

A. Lamp Sensor

The block diagram of the lamp sensor and dimming ballast system are shown in Figure 1. A drawing of the lamp and electrode set up is shown in Figure 2.

The auto-dimming lamp sensor, measures low frequency (~0.1-5Hz) changes caused by human targets below the lamp in the high-frequency (~50kHz) alternating lamp electric fields. The lamp sensor measures the electric fields in front of the lamp with a fully-differential transimpedance front-end amplifier shown in Figure 6(b). The front-end amplifier is fully-differential because both its inputs and outputs are differential. It takes a differential measurement between two electrodes spaced symmetrically about the center of the lamp as depicted in Figure 2. When the electric fields measured by the two electrodes are equal, the differential measurement is balanced and the front-end output voltage is zero Volts or "nulled". Therefore, the front-end can have very high gain without saturating its output in the absence of a detection. This very high gain is necessary to amplify the effects of small imbalances in the capacitive system below the lamp.

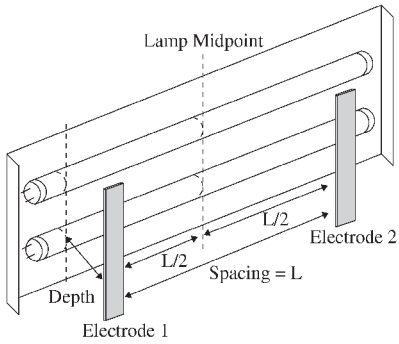

Fig. 2. A diagram of the two-bulb fluorescent lamp and electrodes. The electrodes are spaced symmetrically about the center of the lamp.

For the balanced measurement, the signal source provided by the bulb excitation must consist of two equal signal sources that are symmetric about the center of the lamp. By reversing the connections to one of the bulbs in a two-bulb lamp as shown in Figure 3, each electrode sees "one of each" of the ends of the bulbs. Even if the surface voltage profile for each bulb is nonlinear and asymmetric, the two effective sources behind each electrode will still be equal.

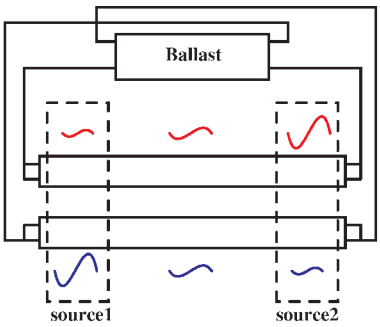

Fig. 3. Reversing the ballast connections to one bulb yields two equal signal sources.

As suggested by the block diagram in Figure 1, the lamp sensor uses synchronous detection to reject stray signals that differ in phase or frequency from the lamp's own excitation signal. The synchronous detector takes a separate measurement of the signal source using another transimpedance amplifier and multiplies it with the main signal using a fully-differential multiplier. Because low-frequency occupant

motion below the lamp modulates the high-frequency signal measured by the front-end, the output of the multiplier contains the low-frequency modulations centered on zero frequency (dc). A low-pass filter (LPF) attenuates the high-frequency residue from the multiplier output to yield the low-frequency signal. The synchronous detector and modulation-demodulation scheme also has the advantage of bypassing potentially overwhelming low-frequency 1/f noise in the front-end amplifier. The signal of interest is amplified at high-frequency (~50kHz) by the front-end and then demodulated by the multiplier. Consequently, the 1/f noise from the front-end amplifier is up-modulated by the multiplier to high-frequency and attenuated by the LPF. This is not unlike chopper stabilization of op-amps for elimination of low-frequency noise, detection in an AM radio or the so-called "lock-in" amplification technique for frequency specificity [8], [9].

In [3], we reported 11ft. of detection range between the lamp and the closest edge of a human target. For a detailed description of the lamp sensor operation please refer to references [1]–[4].

B. Dimming Ballast and Lamp Sensor Interface

A simplified schematic of the dimming ballast is shown in Figure 4. It is designed to operate the lamp between 1.3% and 90% of the full specified power for two T8 32W bulbs. The ballast adjusts the lamp power or brightness by changing its output (excitation) frequency relative to the natural frequency of the LC tank created by L_{res} and C_{res}. The ballast adjusts its excitation frequency from 64.4kHz to 41.7kHz for minimum to maximum lamp power respectively.

The frequency-controlled dimming ballast is designed around the International Rectifier part IR21592 [20]. The ballast consists of a utility-line rectifier and boost-mode power factor corrector (PFC). The STMicroelectronics part L6561 handles the power factor correction and dc-dc up-conversion for the high-voltage dc input to the inverter stage of the ballast. The IR part controls the half-bridge inverter that drives the lamp output stage. The lamp output stage consists of two fluorescent bulbs in series with a "balance transformer", a shunt resonant capacitor and a series resonant inductor. The balance transformer matches the currents in the two bulbs to prevent "winner-take-all" situations in which one bulb strikes before the other, loading down the resonant tank output and preventing the other bulb from striking.

The IR part controls the lamp power as described in the IR21592 datasheet. During dimming, the output stage is effectively an inductance in series with a parallel bulb resistance and resonant capacitance. The current into the output stage is shifted from the half-bridge output voltage somewhere between 0 and -90 degrees during dimming. Zero phase-shift corresponds to maximum lamp power [20].

The IR21592 takes a dim level input voltage (0.5-5Vdc) command and generates a reference signal. The phase between this signal and the gate drive signal is the desired phase between the output current and voltage in order to achieve the proper lamp power. The phase reference signal is compared to the inverter output current and the resulting phase error forces the circuits voltage-controlled oscillator (VCO) to steer the inverter frequency in the right direction. The VCO steers the frequency until the phase error between the output stage current and the reference signal is forced to zero yielding the desired lamp power [20].

The lamp sensor-ballast interface consists of a buffered digital-to-analog converter (DAC) and an optical isolation barrier as indicated in Figure 1. The optical isolator and soft-start circuitry is shown in Figure 5. The optical isolator separates the lamp sensor's ground from the ballast's ground. It also provides a high-voltage safety barrier so that the lamp sensor is not easily damaged by dramatic failures in the ballast. The soft-start is intended to protect the bulbs. It clamps the dimming signal to the range 0.5-5V. It also brings the lower limit down slowly on startup so that the bulbs are struck at a high dimming level and then brought down slowly if the lamp sensor is either off or commanding a low dimming level.

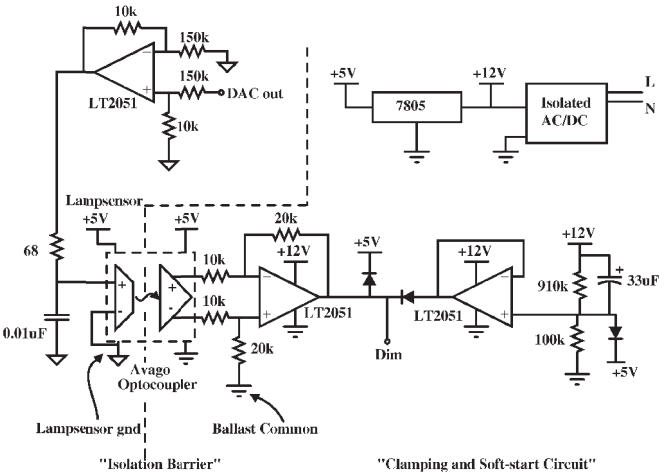

Fig. 5. The optical isolator separates the lamp sensor common potential from the ballast common potential. The soft-start protects the bulb by clamping the dim signal to the range 0.5-5V and brings the dim signal down slowly on startup.

C. Phase Errors

The ballast frequency changes with dimming level. As described in Section II-A, the synchronous detector multiplies the front-end output with the clock amplifier output signal. Therefore, the front-end and clock amplifiers depicted in Figure 1 have to be compensated so that their phases match across the entire frequency range 41.7-64.4kHz. The basic transimpedance-connected op-amp is shown in Figure 6(a). The input capacitance shown corresponds to the lumped capacitances in the lamp system. The circuit to the right of the input capacitance is a transimpedance amplifier because it multiplies the input current by the feedback impedance to produce the output voltage, i.e.

$$v_{out}(\omega_o) = -i_{in}(\omega_o) \times Z_f(\omega_o) \qquad (1)$$

where $Z_f(\omega)$ is the parallel combination of R_f and C_f in Figure 6.

The Bode plot of the closed loop transfer function for the front-end amplifier with a nominal 1pF input capacitance is shown in Figure 7. The two amplifiers were re-compensated so that the entire operating frequency range was within the

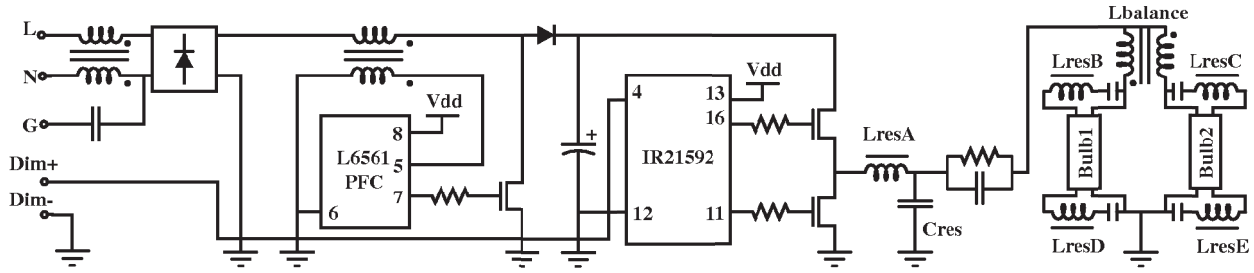

Fig. 4. A simplified schematic of the frequency-controlled dimming ballast. The ballast is designed around the International Rectifier part IR21592 [20].

pass-band or flat part of the frequency response. In the pass-band, the gain of the amplifier shown in Figure 6(a) is $V_{out}/V_{in} = C_{in}/C_f$. The location of the bend or "knee" was chosen to be well below the lowest ballast operating frequency of 41kHz. The knee location is determined by the product of R_f and C_f. To maximize the gain a small C_f was chosen necessitating a large R_f. Because the front-end and clock amplifiers were designed to have very low input bias currents (see references [1]–[4]), R_f may be very large even though it provides the only dc path to the inputs. The final component values are listed in Table II-C.

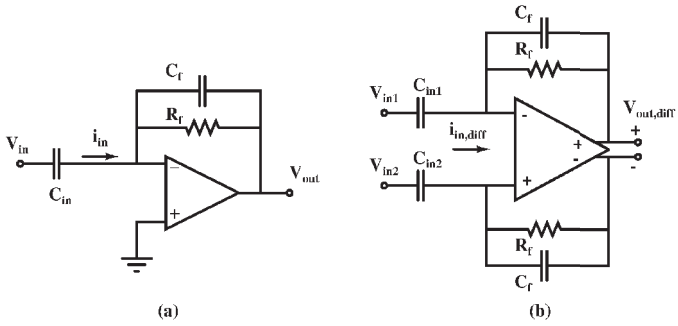

Fig. 6. (a) The circuit to the right of the input capacitor is a transimpedance amplifier. (b) The fully-differential version of a transimpedance amplifier like the one used in the front-end.

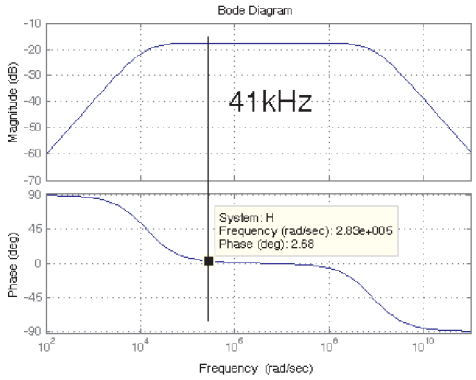

Fig. 7. The new closed loop Bode plots for the front-end amplifier.

The feedback capacitor for the front-end was chosen as a series combination of two 1% tolerance 15pF capacitors. This tolerance is critical to the common-mode rejection (CMR) of the front-end amplifier [1]–[4]. The input bias currents for both the front-end and the clock amplifiers are 10pA maximum [22]. This current develops only 100μV dc across the 10MΩ front-end feedback resistor and less for the clock amplifier.

TABLE I
REVISED FEEDBACK COMPONENT VALUES

Front-end	Clock
$R_f = 10\text{M}\Omega$	$R_f = 220\text{k}\Omega$
$C_f = 7.5\text{pF}$	$C_f = 680\text{pF}$

III. AUTO-DIMMING

This section presents examples of critical system characterizations, a strategy for using those characterizations to design auto-calibration and detection algorithms and a demonstration of the auto-dimming lamp.

Certain performance metrics depend strongly on the height of the lamp and the electrode configuration. For the following discussion of the auto-dimming lamp, the lamp height was 7ft. 4in. above the floor, the electrode depth was 5in. and the electrode spacing was 38in. (see Figure 2).

A. Lamp Sensor SNR across Dimming Levels

The signal-to-noise ratio (SNR) of the lamp sensor system across dimming levels must be determined as this constrains the detection range. To characterize the source, we measured the bulb rms voltage across power levels. It is well known that the bulb current-voltage (I-V) characteristics exhibit positive nominal resistance, but negative incremental resistance for some operating regions [19]. That is,

$$\frac{V}{I} > 0, \text{ but } \frac{dV}{dI} < 0. \tag{2}$$

This means that decreasing the current in the bulb increases the voltage across it. In Figure 8, the rms bulb voltage is plotted against lamp power. Lamp power is calculated as the average power, i.e.

$$\text{P}_{\text{avg}} = (\text{I}_{\text{rms}}\text{V}_{\text{rms}})\cos\phi \tag{3}$$

where ϕ is the phase between the voltage and current. Percent lamp power is defined as the power relative to the maximum bulb power, e.g. 32W. For instance, the lowest lamp power shown, 1.3%, is 0.013×32W = 420mW.

Because lamp current increases monotonically with lamp power, the plot in Figure 8 can also be taken as the bulb's I-V characteristic. The plot shows the familiar negative incremental resistance to the right of the peak in the voltage near 25% lamp power.

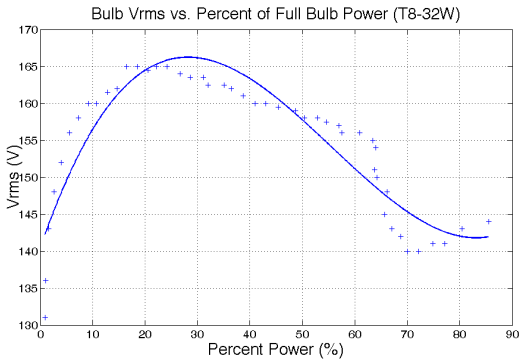

Fig. 8. A curve fit of the rms bulb voltage plotted against lamp power (which increases with rms current) shows the familiar negative incremental resistance of the bulb.

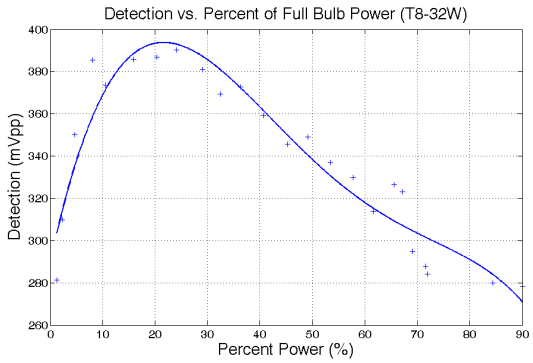

Fig. 9. A curve fit of the detection signal bulb voltage plotted against lamp power.

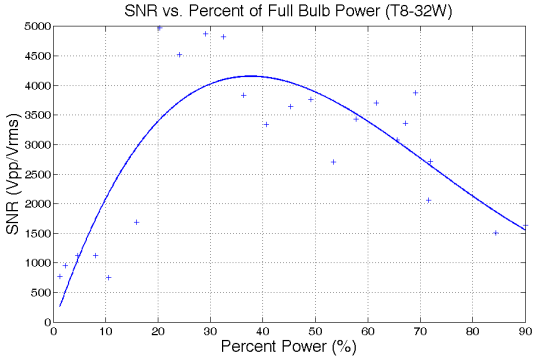

Fig. 10. A curve fit of the SNR plotted against lamp power shows a maximum SNR around 40% power.

To characterize the lamp sensor sensitivity across power levels, we performed an experiment with a hanging lamp sensor to measure the SNR with a 6ft.-tall person under the lamp. Each detection consisted of the target passing through the entire detection field of the lamp sensor yielding a peak-to-peak differential output voltage. Each detection was repeated 5 times for each dimming level. Also, at each dimming level, the windowed noise floor was measured in the time-domain by taking the average ac rms voltage from 15 5-second windows with no target. The noise window length was chosen to match the time it took to fully walk in and out of the detection field. The noise experiments were repeated after the entire first trial and these data were averaged into the first data set in order to minimize the effect of long-term changes in noise levels. The results are presented in Figures 9 and 10. SNR here is defined as the peak-to-peak detection voltage of the human target divided by the ac rms windowed noise voltage in the absence of a detection.

The signal from the lamp sensor increases to follow the bulb I-V characteristic curve so that it shows a maximum around 25% lamp power as shown in Figure 9. However, the sensitivity of the lamp sensor is quantified by the SNR. Because the dominant noise originates in the bulbs and ballast, we might expect it to also vary with the lamp power [1]–[4]. The SNR plot across lamp power levels in Figure 10, shows that the sensitivity actually shows a maximum around 40% lamp power. Such SNR characterizations will be necessary for the designer to determine at which power levels detection is possible. In this experiment, we have shown a situation in which detection is easily possible even at 1.3% (minimum) lamp power.

B. *Characterization of Non-idealities*

Characterization of output voltage offsets and transients when switching between the dim and bright power levels are necessary for the design of auto-calibration and detection algorithms. Despite correcting for phase errors, there are still measurable offsets in the lamp sensor output voltage when switching between dimming levels. These offsets may be due to differences in the voltage profiles of the two bulbs as they are dimmed or due to un-corrected phase errors across dimming frequencies. Figure 11 shows a plot of the output voltage when repeatedly switching between dimming levels. For these particular bulbs and dimming levels (8.1% and 59.4%), the difference between output voltage offsets is about 40mV. The plot shows two important features. First, the offsets are repeatable. That is, when switching back to a dimming level, the output offset will be the same as it was the last time (ignoring drift). Second, there is a consistent transient after each switch, after which the output voltage flattens. The output voltage offsets and transients will determine the baselines for detection in the auto-dimming system. Deviations from the baselines will correspond to detections.

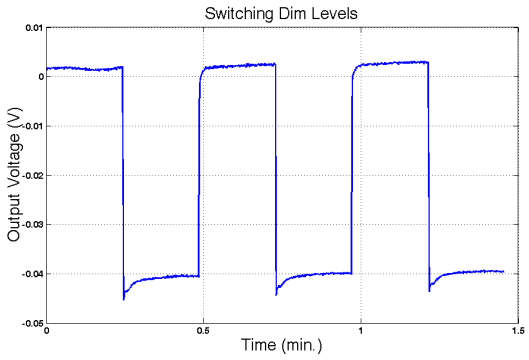

Fig. 11. The lamp sensor output voltage for repeated switching between dimming levels. The two lamp power levels in this example are 8.1%(dim) and 59.4% (bright).

Characterization of output voltage drift ensures repeatable

detection and triggering over time. Figures 12(a) and 12(b) show typical output voltage drift over a period of 60 minutes (long-term) and 3.5 minutes (short-term) respectively for the low lamp power level (8.1%). This data can be used to determine detection bands around the baselines to account for drift. We know from previous experiments, that several mechanisms contribute to drift including sources in the ballast and lamp [1]–[4]. Therefore, we may speculate that warming and cooling of the bulbs yields slowly varying asymmetries between the two signal sources that appear as differential modulations of the measured signal.

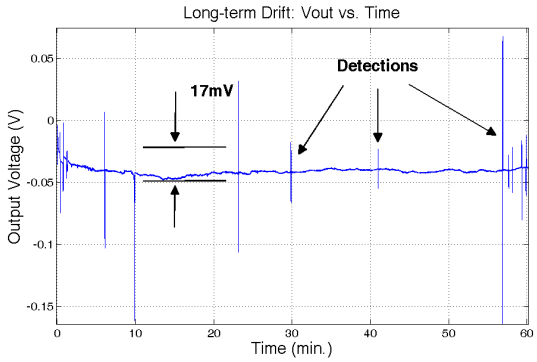

(a) An example long-term drift output voltage plot; 17mV in 60min.

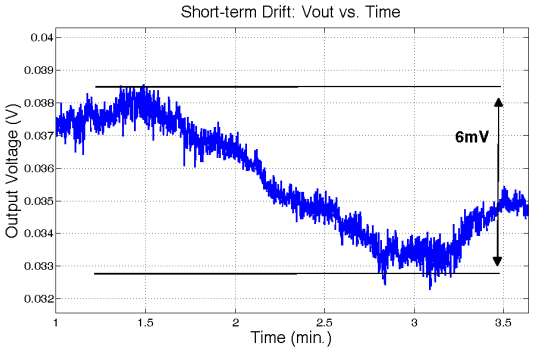

(b) An example short-term drift output voltage plot; 6mV in 3.5min.

Fig. 12. Example Drift Plots for the Lamp Sensor output voltage. Both plots show data taken for a lamp power level of 8.1% (dim).

C. Auto-calibration, Detection and Demonstration

This section shows how to use the characterization of non-idealities to implement an auto-dimming lamp. A screenshot of the output display from the prototype auto-dimming lamp is shown in Figure 13. The screen shot shows two different baselines for when the lamp is bright (59.4% lamp power) and dim (8.1% lamp power) respectively. The baselines are the approximate output voltage in the absence of detection and they are different because of the offsets described above. The two baselines are generated during the auto-calibration which takes place on start-up. The auto-calibration procedure uses the information gained from characterizing the output voltage offsets and transients when switching between dimming levels. The first auto-calibration procedure used was as follows. In the absence of a detection, the lamp is first brightened and the output voltage is measured. The first 80 points (about 6 seconds) are discarded as part of the switching transient depicted in Figure 11 and the remaining 20 points (about 1.5 seconds) are averaged. This average becomes the "bright baseline". The lamp is then dimmed and the procedure is repeated to generate the "dimmed baseline."

The detection algorithm uses the information gained from characterizing the output voltage offsets and drift. The detection algorithm used here was as follows. When the lamp is dimmed, if the output voltage deviates from the dimmed baseline by more than +/-17mV and the deviation persists for more than 200ms, then the lamp is undimmed. When the lamp is brightened, if the output voltage returns to within 10mV of the bright baseline and this persists for more than 200ms, the lamp is dimmed again. The turn-on deviation level (17mV) is greater than the turn-off level (10mV) to discourage multiple transitions. The timing delays serve at least 2 purposes. First, when the target passes under the lamp, it is possible for the output to reach the baseline if the target is symmetrically positioned below the center of the lamp. The delay in this case means that the target must not only be carefully positioned under the lamp but it must stay that way for a little while. Second, the delay adds further protection against multiple transitions when the output voltage is slowly varying.

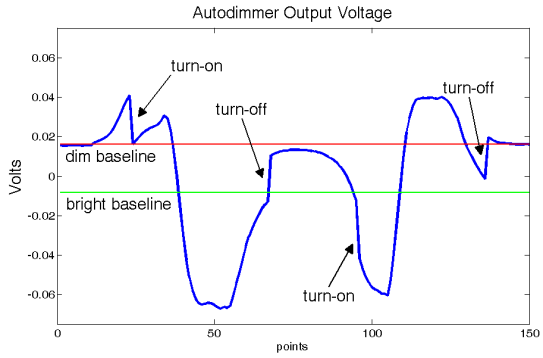

Fig. 13. This plot is a screenshot of the display for the autodimmer. It shows the output voltage varying with the target and the two different baselines taken from the auto-calibration. Jumps in the output voltage can be used to determine when the lamp dimmed or brightened.

More advanced auto-calibration procedures will make the auto-dimmer more robust and more sensitive. One obvious improvement would be to store the transient of the output voltage after switching between dimming levels and use this to adjust the baselines in real-time after switches. More advanced detection algorithms will also improve robustness and sensitivity. For instance, it may be possible to not only constrain the deviation from the baseline, but also the time derivative of the deviation. If the deviation is small, but the combination of the time derivative and the deviation is unlike any drift or noise expected from the lamp sensor, then it may be possible to rule that as a detection. Ultimately, we have the advantage that a "false-positive" detection is not tragic. Therefore, in our detection algorithms we may lean toward avoiding missed detections at the expense of a higher rate of false-positives.

IV. Quasistatic FM Wireless Link

This section details the design and performance of the quasistatic wireless link for communication between adjacent lamps. By communicating with neighboring lamps, a

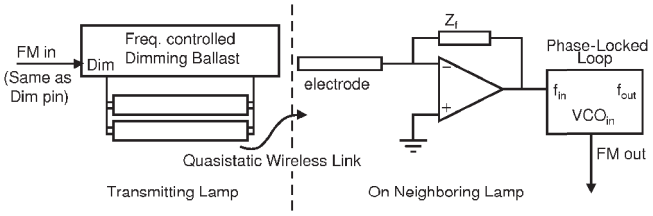

Fig. 14. A notional picture of the quasistatic wireless link system.

detection by a single lamp could, for instance, command its neighboring lamps to turn on as well, and command those lamps to tell their neighboring lamps to turn on. Thus a particular lighting scheme may cause a single lamp, a cluster of lamps, or an entire room of lamps to turn on upon an occupancy detection. Such a network of auto-dimming lamps would self-expand as more lamps were added because the lamps would each passively and wirelessly command their neighboring lamps to turn on.

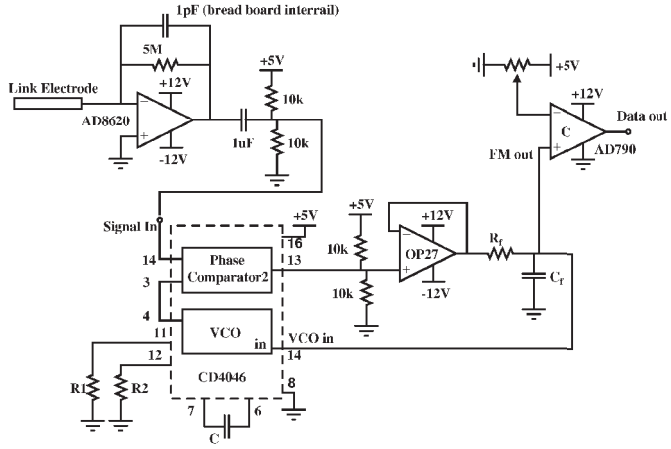

(a) The schematic of the receiver including the PLL.

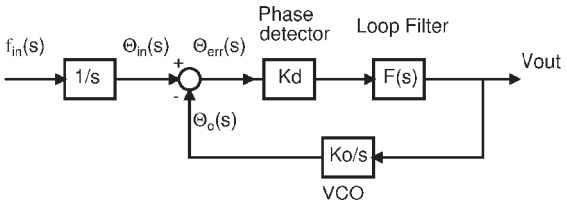

(b) The dynamical block diagram of the PLL from input frequency to output baseband voltage.

Fig. 15. The wireless link receiver and the PLL dynamical block diagram.

The wireless link is quasistatic because it does not use propagating electromagnetic waves for transmission. It uses the same electric fields from the lamp that the lamp sensor uses to detect occupants. The electric fields vary slowly enough (~50kHz) that the associated wavelengths for propagating waves at these frequencies are extremely long compared to the length-scales of the lamp sensor system. It is a frequency-modulated (FM) link because it transmits information by adding small frequency deviations or modulations (~+/-500Hz) to the high-frequency (~50kHz) ballast signal.

The notional drawing of the wireless link is shown in Figure 14. The baseband or frequency modulating signal is input as small voltage deviations on the dim pin of the dimming ballast. A single-ended transimpedance amplifier, similar to those used in the lamp sensor, measures the frequency-modulated electric fields at the receiving lamp's electrode. The signal that is detected is input to a phase-locked loop (PLL). When the PLL locks onto the input signal, the VCO input is a dc voltage which corresponds to the input signal frequency. When the input signal frequency changes, the PLL tracks those changes and the VCO input varies correspondingly. The baseband or modulating signal is then available at the VCO input, which we call the output of the PLL here.

A. PLL Design

The full schematic of the receiver including the PLL is shown in Figure 15(a). The dynamical block diagram is shown in Figure 15(b). This PLL uses a phase detector and VCO from the Fairchild Semiconductor part CD4046 [23]. The gain term, K_d, is determined by the phase detector output voltage function for the connection shown in Figure 15(a) [23]:

$$V_{pd} = 2.5\text{V} + \frac{\Theta_{\text{err}}}{2\pi} \times 2.5\text{V} \tag{4}$$

$$K_d = \frac{dV_{pd}}{d\Theta_{\text{err}}} = \frac{2.5}{2\pi}\frac{\text{V}}{\text{rad}}. \tag{5}$$

The gain term K_o is the VCO gain from input voltage to output frequency. The external components R_1, R_2, and C determine K_o and the VCO center frequency, f_o. Using the CD4046 datasheet, these values were chosen to yield

$$f_o = 57.4\text{kHz} \tag{6}$$

and

$$K_o = 14.4\frac{\text{krps}}{\text{V}}. \tag{7}$$

Smaller modulations of the ballast frequency yield less visible flicker in the light. The gain K_o could be as low as necessary for those modulations to fill the input dynamic range of the VCO (0.5-4.5V). However, such a small K_o will lead to a small gain-bandwidth product (GBW), K_oK_d, of the integrator in the loop. The transfer function from f_{in} to $\text{VCO}_{in} = V_{out}$ with a passive RC low-pass loop filter is

$$\frac{V_{out}(s)}{f_{in}(s)} = \frac{K_d/R_fC_f}{s^2 + s/R_fC_f + K_dK_o/R_fC_f}. \tag{8}$$

Recognizing K_dK_o as the GBW and the entire denominator of eqn. (8) as the characteristic polynomial of a second-order system,

$$p(s) = s^2 + 2\zeta\omega_n s + \omega_n^2 \tag{9}$$

$$\omega_n = \sqrt{\frac{\text{GBW}}{R_fC_f}} \tag{10}$$

$$\zeta = \frac{1}{2}\sqrt{\frac{1}{R_fC_f \times \text{GBW}}} \tag{11}$$

where ω_n is the natural frequency and ζ is the closed-loop damping ratio. From [6], the bandwidth of this second-order system is

$$\omega_h = \omega_n\sqrt{1 - 1\zeta^2 + \sqrt{2 - 4\zeta^2 + 4\zeta^4}}, \tag{12}$$

which is always greater than ω_n so that the closed-loop bandwidth, ω_h, increases with the GBW, K_oK_d, of the loop.

A small K_o will also limit the frequency range which the PLL can lock onto (lock range). Instead, choosing a larger value of K_o increases the GBW of the loop and the lock range. As a consequence, small frequency modulations only produce small deviations of the VCO input voltage, so the threshold of the comparator at the output of the system must be trimmed to be centered on those small deviations. A tradeoff between noticeable flicker, PLL noise and PLL bandwidth falls out of this discussion. If, for instance, we wanted to reduce the change in dimming levels to reduce the noticeable light flicker, we could reduce the change in ballast frequencies or modulation depth but this would also reduce the VCO input deviations. At some point the VCO input devations will fall below the VCO input-referred noise. To increase the deviations we only have two choices: increase the flicker in the lamp, or decrease the gain K_o of the VCO so that the reduced changes in frequency fill more of the VCO input range. However, from above, we know that decreasing K_o decreases the closed-loop bandwidth of the PLL. We can write a fundamental expression relating these trade-offs in the PLL described here.

$$\text{GBW} = K_oK_d \tag{13}$$

$$K_o\Delta\text{VCO}_{in} = \Delta f \tag{14}$$

$$\boxed{\frac{\text{GBW}}{K_d} = \frac{\Delta f}{\Delta\text{VCO}_{in}}} \tag{15}$$

To attenuate the high-frequency components of the phase-detector output while maintaining reasonable loop stability, the LPF 3db frequency was chosen to match the cross-over frequency of the integrator in the feedback loop shown in Figure 15(b). That is,

$$\omega_{3db} = \frac{1}{R_fC_f} = K_oK_d = \text{GBW} \tag{16}$$

This strategy leads to about 45° of phase margin. The values for R_f and C_f in Figure 15(a) are

$$R_f = 2\text{k}\Omega \tag{17}$$

$$C_f = 0.1\mu\text{F}. \tag{18}$$

The Bode plots of the loop transfer function and the closed-loop transfer function are shown in Figure 17(a) and 17(b). They show 49° of phase margin and a closed-loop bandwidth of 6.9krps or 1.1kHz. In this PLL, because the loop filter is a first-order RC LPF, the average value of the phase detector output is the VCO input. Therefore, the steady-state phase error, $\Theta_e(t)$ and the corresponding phase detector output voltage varies with the steady-state VCO output frequency. Because the VCO input range is 0.5-4.5V, the extremes of the output frequency lock range in eqn. (4) correspond to a phase error range of about +/-5 rad. Therefore, for all frequencies within the lock range, the phase error is never more than +/-2π rad. and this PLL is not vulnerable to skipping cycles during step transients or "cycle slipping" [7].

The simulated and measured step response from the input signal to the VCO input voltage are shown in Figures 16(a)

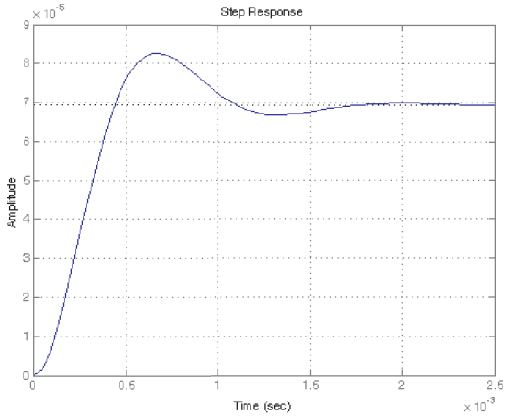

(a) Simulated step response from the input signal to the VCO input in the PLL.

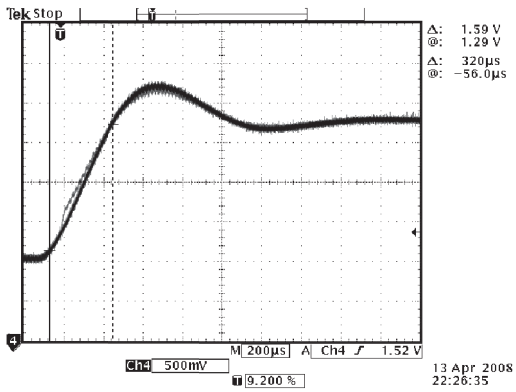

(b) Measured step response using an FM-capable Agilent signal generator.

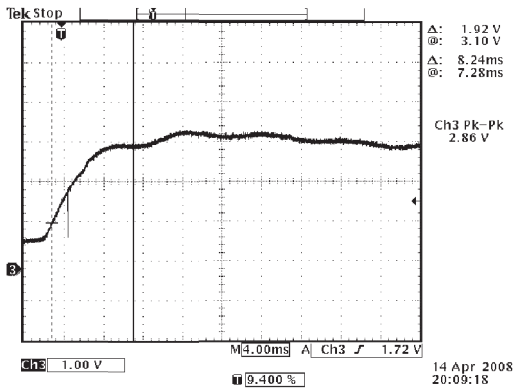

(c) The PLL step response with the lamp in the system shows a slower rise time.

Fig. 16. Simulated and measured step responses of the PLL.

and 16(b). The measured step response was generated with an FM-capable Agilent signal generator driving the input signal to the PLL. The measured peak overshoot in the step response is about 1.2 which for a second-order system corresponds to 45° of phase margin as predicted above [6]. The rise time of the measured step response is about 320μs which for a second order system corresponds to a bandwidth of

$$\text{BW} \approx \frac{2.2}{t_r} = 6.9\text{krps} = 1.1\text{kHz}, \tag{19}$$

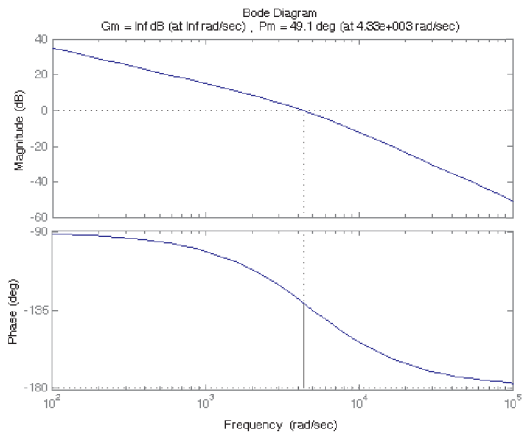

(a) Bode plot of the loop transfer function of the PLL.

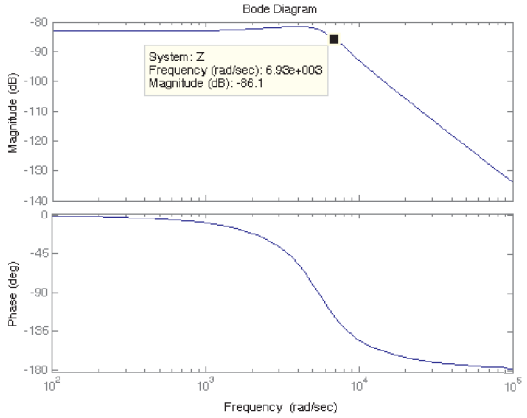

(b) Bode plot of the closed-loop transfer function of the PLL.

Fig. 17. Simulated PLL dynamics.

also as predicted above [6].

Finally, Figure 16(c) shows the step response using a transmitting lamp in place of the FM signal generator. The lamp frequency in this example is modulated between 56.6kHz and 55.6kHz by modulating the lamp power between 36.3% and 45.3%. The rise time in Figure 16(c) is about 8.24ms. This corresponds to a bandwidth of

$$\mathrm{BW} = 43\mathrm{Hz} \tag{20}$$

The PLL design is sufficient because the lamp is clearly the limiting factor in the signal bandwidth.

B. Wireless Link Demonstration and Range

Figures 18(a) and 18(b) depict bitstreams transmitted with the wireless link. For the system demonstrated here, the lamp power is modulated between 36.3% corresponding to $f = 56.6\mathrm{kHz}$ and 45.3% corresponding to $f = 55.6\mathrm{kHz}$. These frequency deviations result in peak-to-peak VCO input voltage deviations of 700mV centered on 1.9V near the center of the VCO input voltage range. The output data is available at the comparator output in Figure 15(a).

The maximum distance between adjacent lamps for proper wireless link operation affects lighting network design. To determine the link range, one lamp was used to transmit data while a receiving electrode was moved away until there

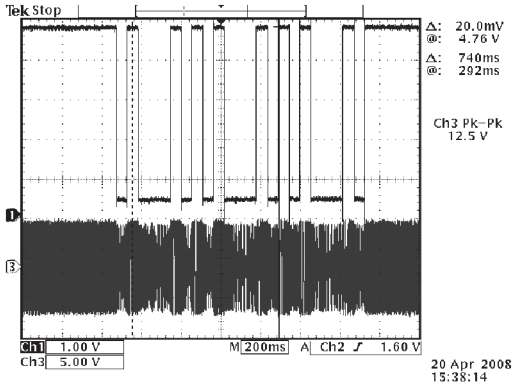

(a) A 3-byte long (24-bit) packet transmitted over the wireless link and demodulated by the PLL (top) and frequency modulations are evident in the output of the receiver front-end before demodulation (bottom).

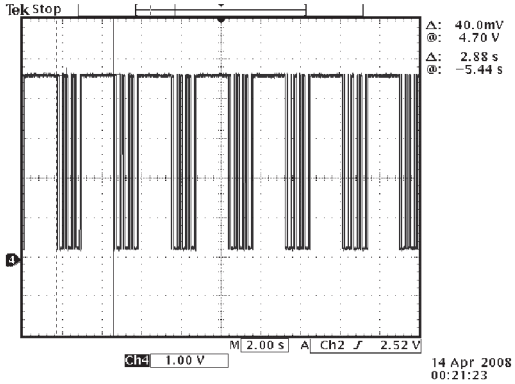

(b) A sequence of 7 3-byte long packets.

Fig. 18. Wireless link demonstration

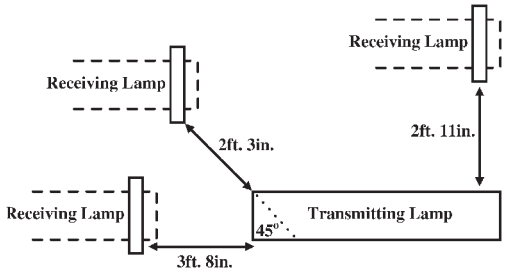

Fig. 19. The measured wireless link range between the closest edge of the transmitting lamp and the receiving lamp's electrode.

were visible bit errors in the received packets. A dummy lamp under the receiving electrode modeled the capacitive structure in a realistic system. As depicted in Figure 19, the adjacent lamp end-to-end maximum distance was 3ft. 8in. The range between two parallel lamps (broad-side range) was 2ft. 11in. The 45° range between the nearest corner of the transmitting lamp and the nearest corner of the receiving electrode was 2ft. 3in.

V. Conclusion

The autonomous demand-side energy management system for control of lighting energy consumption was motivated by energy conservation and efficiency efforts and was supported by previous work in Section I. We have demonstrated the building blocks of a new autonomous demand-side energy management system to efficiently control lighting

energy consumption with very fine-grain measurements of occupancy. The demand-side energy management system uses fluorescent lamp sensors interfaced with dimming ballasts. The system in this paper demonstrated high-sensitivity (SNR) even with the lamps dimmed to only 1.3% of the maximum lamp power. This implies a 98.7% power savings for each lamp while it is dimmed. Furthermore, we were able to network the lamps using a wireless link that re-used the ballast action and the same electric field measurement techniques that the lamp sensor used. Finally, we have provided guidelines for the system designer to determine critical system characterizations in order to implement a network of auto-dimming lamps. For the system described in this paper, the most important system specifications are shown in Table V. In the table, the lamp height and electrode configuration are also recorded. Refer to Figure 2 for details of the electrode configuration.

TABLE II
DEMONSTRATED AUTO-DIMMING LAMP PERFORMANCE

Spec	Value
Lamp Height	7ft. 4in.
Electrode Depth	5in.
Electrode Spacing	38in.
Min. Lamp Power for Detection	1.3% (420mW)
Min. SNR	750Vp-p/Vrms
Max. SNR	4200Vp-p/Vrms
Min. SNR Lamp Power	1.3%
Max. SNR Lamp Power	38%
Typ. 60 min. drift	17mV
Typ. 3.5 min. drift	6mV
Link end-to-end range	3ft. 8in.
Link 45° range	2ft. 3in.
Link Broad-side range	2ft. 11in.
Link Bandwidth	43Hz

ACKNOWLEDGMENTS

This research was funded by the MIT Energy Initiative, the U.S. Department of Justice, the NEMOmetrics corporation, the Grainger Foundation, and by a grant from the National Science Foundation.

REFERENCES

[1] John J. Cooley, Al-Thaddeus Avestruz, Steven B. Leeb "Proximity Detection and Ranging using a Modified Fluorescent Lamp for Security Applications". *Carnahan Conferences Security Technology, Proceedings 2006 40th Annual IEEE International*, 2006.

[2] John Cooley, Al-Thaddeus Avestruz, Steven B. Leeb "A Modified Fluorescent Lamp for Discreet Biometric Surveillance". *Conference Homeland Security, Proceedings 2007 IEEE*, 2007.

[3] John J. Cooley, *Capacitive Sensing with a Fluorescent Lamp*, MIT Master of Engineering Theses, 2007.

[4] Cooley, J.J., Avestruz, A.-T., Leeb, S.B., Norford, L.K., "A Fluorescent Lamp with Integral Proximity Sensor for Building Energy Management". *Power Electronics Specialists Conference, 2007. PESC 2007. IEEE*, 17-21 June 2007 Page(s):1157 - 1163

[5] James T. Dakin. "Nonequilibrium Lighting Plasmas". *IEEE Transactions on Plasma Science*, Vol. 19, No. 6, 1991.

[6] James K. Roberge. *Operational Amplifiers: Theory and Practice John Wiley and Sons, Inc., New York, London, Sydney, Toronto, 1975.*

[7] Ronald E. Best. *Phase-Locked Loops: Design, Simulation, And Applications McGraw-Hill, New York 1999.*

[8] Christian C. Enz, Gabor C. Temez. "Circuit Techniques for Reducing the Effects of Op-Amp Imperfections: Autozeroing, Correlated Double Sampling, and Chopper Stabilization." *Proceedings of the IEEE, Vol. 84, No. 11, November 1996.*

[9] Andrew E. Moe, Steve R. Marx, Iqbal Bhinderwala, and Denise M. Wilson "A Miniaturized Lock-in Amplifier Design Suitable for Impedance Measurement in Cells" *Sensors, 2004. Proceedings of IEEE* 24-27 Oct. 2004 Page(s):215 - 218 vol.1

[10] Yao-Jung Wen; Granderson, J.; Agogino, A.M. *Towards Embedded Wireless-Networked Intelligent Daylighting Systems for Commercial Buildings Sensor Networks, Ubiquitous, and Trustworthy Computing, 2006. IEEE International Conference on* Volume 1, 05-07 June 2006 Page(s):326 - 331

[11] Advanced monitoring technologies for the evaluation of demand-side management programs De Almeida, A.T.; Vine, E.L.; Power Systems, IEEE Transactions on Volume 9, Issue 3, Aug. 1994 Page(s):1691 - 1697 Digital Object Identifier 10.1109/59.336086

[12] Right power, right price [enterprise energy management systems] Forth, B.; Tobin, T.; Computer Applications in Power, IEEE Volume 15, Issue 2, April 2002 Page(s):22 - 27

[13] Sensor/actuator networks supporting agents for distributed energy management Taylor, K.; Ward, J.; Gerasimov, V.; James, G.; Local Computer Networks, 2004. 29th Annual IEEE International Conference on 16-18 Nov. 2004 Page(s):463 - 470

[14] "Convservation at the Touch of a Button." Matt Richtel New York Times Article, November 7, 2007.

[15] U.S. Department of Energy "Annual Energy Outlook." Energy Information Administration, February 2007.

[16] U.S. Environmental Protection Agency "Energy Management and Conservation Program: Fiscal Year 2005 Annual Report"

[17] U.S. Environmental Protection Agency EPA Recognizes Energy Star Winners for Outstanding Energy Efficiency, March 21. *EPA Newsroom.* 2007.

[18] U.S. Environmental Protection Agency "Energy Star and Other Climate Protection Partnerships: 2006 Annual Report" 2006.

[19] Shmuel Ben-Yaakov, Moshe Shvartas, and Stanislav Glozman. "Statics and Dynamics of Fluorescent Lamps Operating at High Frequency: Modeling and Simulation". *IEEE Transactions On Industry Applications, Vol. 38, No. 6*, Dec. 2002

[20] IR21592 Data Sheet International Rectifier, El Segnudo, CA 90245 9/7/2005

[21] L6561 Data Sheet STMicroelectonics Group of Companies, 2004

[22] AD8620 Datasheet Analog Devices, Norwood, MA, 2004

[23] CD4046 Datasheet Fairchild Semiconductor, Octobe 1987

[24] HCPL-7800A-000E Datasheet Avago Technologies, 2007