An LED-to-LED Visible Light Communication System with Software-Based Synchronization

Stefan Schmid*†, Giorgio Corbellini*, Stefan Mangold*, Thomas R. Gross†

*Disney Research
8092 Zurich, Switzerland

†Department of Computer Science
ETH Zurich, Switzerland

***Abstract*—An LED can emit and receive light and provides therefore a simple building block for a Visible Light Communication (VLC) system. We describe a microcontroller-based system and report on its effectiveness in a testbed. The key idea is to use a microcontroller to provide synchronization so that the receiver can lock to the transmitted signal in a fast and efficient way. Further we propose a combined light emission and light measurement approach to receive with an LED while emitting light. The proposed system enables new entertaining applications by creating the illusion (for a human observer) that both transmitting and receiving LEDs are always switched on.**

I. INTRODUCTION

Light communication has several advantages: Light communication is visible (in contrast to invisible radio communication), so it is easy to determine who can listen to (or receive) a message. Furthermore, light communication does not use radio waves, and there are environments or communities that may value this aspect. A side effect is that light communication does not require part of the (limited) radio spectrum and can therefore be seen as a suitable extension in bandwidth-limited scenarios. And (visible) light is present in many places, so there is the opportunity to combine light communication with lighting design to let Visible Light Communication (VLC) coexist with (or even benefit from) the lighting setup present in many offices, homes, or institutions.

The VLC principle is a relatively new approach for optical free space applications. However, it has been so far considered mainly for Internet access or home networks, but more creative use cases are feasible. For example, VLC can be employed for toy cars or may allow a magic wand to control light effects on a dress [1],[2].

LEDs provide an almost ideal platform for VLC. An LED can emit and receive light at the same time (with multiplexing). In the present paper we present an LED-to-LED communication system for Visible Light Communication (VLC). Such a system can modulate light intensity with high frequencies so that the human eye is not affected by the light communication; the eye always perceives a continuous light even though optical messages are transmitted.

The system has a low complexity and uses LEDs to achieve bidirectional communication avoiding the use of photodetectors. Bidirectional communication is obtained by temporal separation of transmitter and receiver signals. To achieve bidirectional communication using a slotted carrier sense multiple access (CSMA), devices must be synchronous. In this paper, we describe a synchronization model where the receiver is able to lock to the transmitted signal in a fast and efficient way. The system has been realized in a testbed that allows us to gain novel insights. The system uses single LEDs for transmission and reception using an On-Off Keying modulation; the result is a low-power system for low data rate applications.

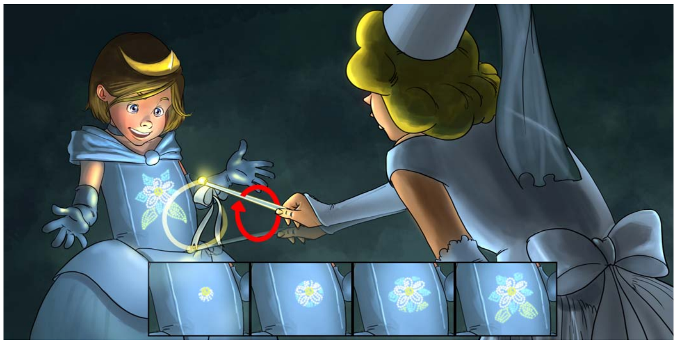

Fig. 1. Typical VLC application scenario for consumer products. LEDs are used as light emitters and receivers. Toy magic wands are used to communicate with LEDs embedded in fashion.

The rest of the paper is organized as follows. In Section II we present the technological background of the system. In Section III we provide a detailed description of the testbed. Performance issues encountered in the testbed are discussed in Section IV. Section V summarizes briefly related work, and we present our conclusions in Section VI.

II. TECHNOLOGY

A simple, off-the-shelf LED can be used not only as a light emitter but can also replace photodetectors and fulfill their task as a light receiver [3]. We consider that each device is equipped with one LED that is used for both transmission and reception. We now focus on the description of the receiver design, assuming that the sender transmits a periodic idle pattern, an optical signal with period 500μs consisting of no-light (corresponding to two symbols '0') and light (corresponding to two symbols '1') slots. The resulting signal frequency is 1 kHz, and this value is chosen because most humans are sensitive to light flickering below 1 kHz. To prevent flickering and brightness changes due to the transmission of other messages than the idle pattern, such as data messages, we use Manchester Coding [4] to distribute the occurrences of '1' and '0' more evenly; the human observer will perceive an always-on optical signal.

A. LEDs as a Substitute for Photodetectors

The following setup can be used to measure with an LED the amount of received light: The incoming photons generate a small photo-current inside the LED. Since an LED is not optimized to receive light, the induced photo-current is very small. To make it easier to measure this slight current, the LED can be set up in reverse bias. In this setting the LED acts as a small capacitor that is charged by putting the LED's cathode to 5V. In this charged state, the photo-current generated by the photons accelerates the discharging of the capacitor. If a analog-to-digital-converter (ADC) is available, the voltage over the LED can be measured after a constant time period. Low ADC-values (low voltage) indicate a fast discharge (high photo-current), and high ADC-values (high voltage) mean that there was no light-source present. If there is no ADC available, a digital input can be used to decide if the voltage, measured after the sampling time, is close to logic HIGH or LOW. This solution is limited to short distances between the light-source and the receiver, and the results depend also heavily on the intensity of the transmitter. The LED might not receive enough light to ever drop below the LOW voltage-threshold, and in this case it is impossible to distinguish transmitted light from ambient light. The ADC, on the other hand, allows the receiver to set its own threshold and to detect finer changes in the measurements.

In summary, by using an LED and an adaptive threshold, it is possible to decide if there was light or no-light (or even partial light) during the past measurement slot. This fact can be exploited to receive and decode a message transmitted by another LED. Furthermore, we used On-Off-Keying [4] on the physical layer to encode and decode a message.

B. Sampling and Synchronization

The Manchester symbols can be decoded at the receiver's end by sampling (measuring the incoming light) with a constant period equal to the transmission period. At the beginning of each sampling interval, the LED (capacitor) is charged to 5 V for 1 μs. To start the discharging process the cathode of the receiver LED is configured as an input. At the end of each sampling interval, the remaining voltage can be measured and stored for later processing. To be able to correctly detect which symbol is transmitted, the measurement intervals must be synchronized to the sender's signal period. The synchronization makes it possible to demodulate a message without oversampling the received signal and therefore increases the overall channel capacity. The synchronization method is further described in Section III.

To create the illusion of an always-on LED also at the receiver side, it is possible to alternate measuring slots and light emitting slots following a specific pattern. If not every time slot is used to measure the incoming light, in fact, it is possible to use the non-measuring slots to emit light. Fig. 2 shows two measuring slots always followed by two light emitting slots. If both the transmitter and the receiver follow the same pattern, called *idle pattern*, the signal generated this way is still 1 kHz (if measurement slots last 500 μs) and both LEDs are still not perceived as flickering.

The multiplexing of the LED introduces also new challenges. Measurements show (see Section IV) that the result of a light measurement depends on what happened in the slot before, even if the amount of received light is the same. In this setup, the first measurement follows a light emission slot (Fig. 2), and the second measurement is preceded by a measurement slot. Exact and trustworthy measurement values are the key for a successful synchronization and it is therefore necessary to create the same preconditions for every measurement. Section IV reports on how this issue is solved in our implementation.

III. IMPLEMENTATION

An off-the-shelf microcontroller often provides only one ADC unit and therefore only one ADC conversion at the time is possible. Moreover, another design constraint is the limited computation power of simple devices. As a result, it is not possible to process more than one signal during the time occurring between two measurement slots. This problem can be solved by multiplexing multiple LEDs over time. An LED is measuring the light for a fixed number of slots. If it recognizes a message, the microcontroller continues measuring using the same LED. Otherwise, if nothing can be decoded and the fixed number of measuring slots is reached, the microcontroller switches to another LED. Since one measuring slot is 500 μs, the measuring loop over all LEDs is perceived by a human as all LEDs are measuring at the same time. Table I shows the key specifications of the hardware used for the implementations. Fig. 3 summarizes the building blocks of the implemented Transmitter and Receiver. The different building blocks are addressed in the following subsections.

A. Light Measurement

At the beginning of each measurement slot, the anode is set to ground and the cathode to 5 V. After 1 μs, the cathode on the microcontroller is switched to input and the pull-up resistor is disabled. The LED, now acting as a capacitor, starts to discharge itself while incoming photons are accelerating its discharge (see Fig. 3, PHY block of the Receiver). At this point a hardware timer is reset to trigger the next measurement slot. The time between the charge and the actual measurement at

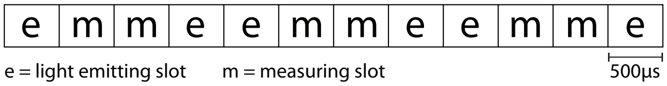

Fig. 2. Alternating light emitting slots with light measurements slots creates the illusion of an always-on LED. The human eye is not able to recognize the fast switching between the two slot types.

TABLE I
KEY PARAMETERS OF THE IMPLEMENTED VLC SYSTEM.

Property	Value
Microcontroller clock	16 MHz
Peak voltage	$5V$
Signal period	$500\mu s$
ADC resolution	10 bit
ADC clock	1 MHz
ADC buffer size	32 samples
LED size	5 mm
LED color	RGB-LED (using mostly red for tx)

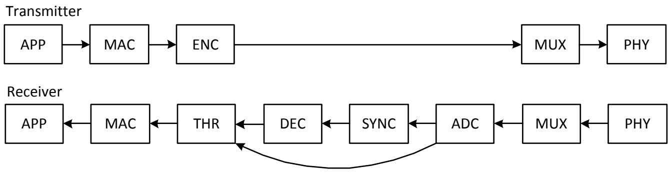

Fig. 3. Block diagram of the receiver and the transmitter. A encoder/decoder (ENC/DEC) Manchester coding is used. A multiplexer (MUX) switches between receive light and light emission modes, and an adaptive threshold (THR) is used to decode '0' and '1' from the ADC values.

the end of the slot can be used to process the measurement of the previous slot. The value is retrieved from the ADC registers and is placed in a buffer containing the last measurement values. A timer interrupt then triggers a handling routine, which starts the analog-to-digital conversion to measure the cathode's potential at the end of the measurement slot. The ADC unit triggers another interrupt upon finishing the conversion which again initiates the next measurement slot (Fig. 3, ADC of the Receiver).

B. Combined Light Emission and Light Measurement

As described in Section II, it is possible to emit light between two measurement slots to emulate an always switched-on LED. To avoid flickering, it is important to not go below 1 kHz. A pattern of two measurement slots followed by two light emitting slots, as shown in Fig. 2, provides the possibility to synchronize in an efficient way and also allows to decode the same Manchester encoded signal without loosing channel capacity (Fig. 3, LED Mux of the Receiver).

C. Processing and Decoding Threshold

After retrieving a value from the ADC registers, it is stored in the buffer of the last received values. To decode a signal out of the stored ADC values, the system must decide whether a certain value relates to a logic HIGH or a logic LOW (Fig. 3, Threshold block of the Receiver). This can be done by using a detection threshold. A static threshold can be used, based on experiments that depend heavily on the current setting including the light intensity (and distance) and the ambient light. However, to make the system work independently of the environment, an adaptive threshold must be introduced. We compute the threshold as the mean of all values in the ADC-values buffer. The mean can be computed using a fast and incremental algorithm that provides a good threshold if the distribution of the light and no-light slots is uniform. However, in the implemented system both the transmitter and the receiver follow the idle pattern (Fig. 2), as a result, a threshold equal to the average of the last measured ADC values cannot be used directly. If they are synchronized in fact, the threshold computed from the mean of the latest ADC results is always very close to the value of a full light slot as the measurements are taken only during the light slots (or part of the light slots). During all the dark slots (only ambient light) of the transmitted signal, the LED emitted light and could not measure anything. Hence, the resulting mean does not represent a good threshold. Since the mean value represents only full light values and the ambient light can be approximated as no light, halving the computed mean value provides a good threshold value for the decider. Further, during the reception of a data frame, the threshold adaption is stopped to preserve a good threshold for the message decoding and again started afterwards.

D. Synchronization

Fig. 4 shows a snapshot of the synchronization process on the oscilloscope. The signal at (4) is the idle pattern transmitted by the sending LED. The signal relates to a binary '0110 0110', two $500\,\mu$s light slots, followed by two dark slots (ambient light). To decode a message correctly, the measurement slots indicated by the peaks at (3) must be in phase with the signal at (4). In the figure, we see that the measurement slot starts around 60μs too early (time between the two vertical bars, indicated as Δt at the bottom of the figure). The synchronization works as follows: As long as the two measurement slots are not reporting a similar value, the signals are not in phase. By comparing the two values, it is even possible to detect in which direction the measurement slot must be shifted to reach a synchronized state. If there is less light detected in the first slot, i.e., there is a higher measured voltage, we are sampling too early and we must shift to the right, that is, measure later (as shown in the figure). If the second voltage value is higher instead, we measured too late and must shift to the left. Shifting adjustment is repeated until a stable state is reached. In the figure, the last measuring slot for the signal at (3) shows an increased duration. This indicates that the synchronization algorithm detected that the phase was not synchronized for the last two measurements slots (indicated by the peak of (1)). By increasing the slot duration time, the measurement shifts to the right and is now synchronized with the transmitted signal at (4). The peaks shown at (2) mark the detection of a light value above the threshold resulting in a decoded '1' (Fig. 3, Sync block of the Receiver).

E. Message Format

Fig. 5 illustrates the message format provided by the MAC block that is sent to the Manchester encoder. The message format includes 5 fields: The Start Frame Delimiter (SFD), the message length, the MAC Header (MHR), the data, and the Cyclic Redundancy Check (CRC). The SFD indicates

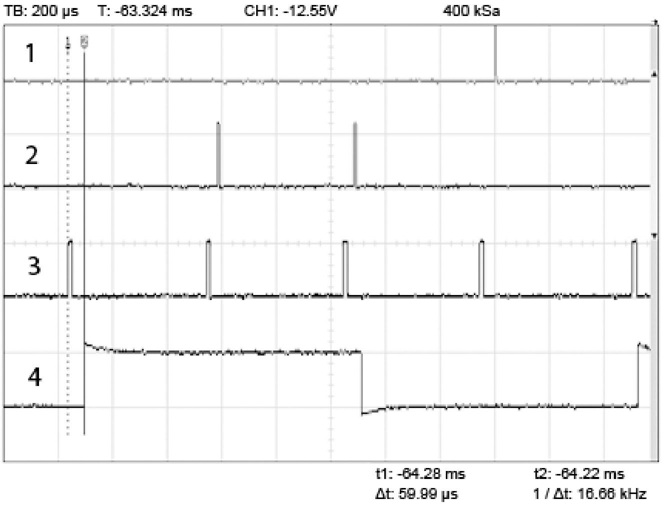

Fig. 4. Snapshot of the synchronization procedure. The rows (1), (2), and (3) belong to the receiver, row (4) presents the optical signal transmitted by the sender. The receiver adjusts the synchronization phase in (1). The peaks in (2) correspond to the detection of symbols '1' (no peaks correspond to a '0'). Slot starting points are depicted in (3).

the start of a new message and is encoded with a sequence of 8 bits. The SFD indicates the beginning of a new data frame and its correct detection is therefore fundamental for a correct data reception. It must be clearly distinguishable from the idle pattern that is continuously transmitted and used for synchronization. Thanks to the efficiency of the synchronization process and the precision of the computed threshold, the implemented SFD sequence is '1010 1011', which has only 1 bit difference (2 encoded with Manchester) from the idle pattern. The rest of the message is similar to the one introduced in the 802.15.7 standard [5]. The message length field has 8 bits and is followed by the MHR. The MHR consists of three bytes, one for the MAC Frame Control (MFC), one for the destination address, and one for the source address. The MFC consists of three bits for the frame type, one bit for a "more data pending" flag and an ACK request flag. The remaining three bits are unused. After the MHR, the data payload is sent. A MAC footer containing the 16-bit CRC checksum over the MHR and data payload is appended. The ACK messages use a MAC layer format without payload (data size equal to zero bytes), simply acknowledging a successful reception (Fig. 3, MAC block of the Receiver).

F. Manchester Encoder and Decoder

Bits are encoded using a Manchester encoder to avoid light flickering and giving to the human observer the perception of continuous light emission independently of the type of message that is being sent. This results in an uniformly distributed zero-one pattern but doubles the needed message symbols (Fig. 3, Manchester encoder, Transmitter).

IV. EVALUATION

This section presents the performance of the implemented VLC system. We used an ATmega328P [6] microcontroller equipped with RGB-LEDs. For the detailed specifications see Table I. Further, the evaluation was executed using the combined light emission and light measurement approach described in Section III.

8 bit SFD	8 bit length	24 bit MAC frame header	0-2040 bit data payload	16 bit CRC

Fig. 5. MAC frame structure - the start frame delimiter (SFD) indicates the beginning of a new frame, the length field defines the data payload size, the MAC frame header contains control flags, and the CRC enables frame integrity checks.

A. Measurement Correction

The synchronization process described in Section II-B relies on the precision of voltage measurements at the end of each measurement slot. The current implementation of the system uses an ADC with a 10-bit precision to represent voltage values in the range [0-5V]. The photon emission of an LED depends also on its thermal state. Therefore, the photo-current generated by incoming photons is also influenced by the current temperature of the LED. Since the size of the photo-current is responsible for the resulting voltage at the end of a measurement slot, the detected ADC values depend on the LED's current thermal state. An idle pattern composed of light and dark slots creates different thermal preconditions for a light measurement. For example, the first measured light slot after a dark slot is a "cold" measurement and results in a different ADC value with respect to the following measurement, even if the light received is constant. Using the combined light emission and light measurement setup adds even more thermal factors, since using the LED for light emission will also heat it up. The thermal effect only accounts for a small part of the difference. After a light slot (after switching off the LED), the microcontroller's ground changes for a short amount of time due to the additional available current. This also changes the measured value since the ADC evaluates the wanted voltage between ground and the reference voltage. Fig. 6 shows these effects on the measured ADC values. The plot is obtained using the measuring schema depicted in Fig. 2 and the corresponding '0110 0110' pattern (always light during measurement slots) perfectly synchronized to the measurement slots. The plot shows the difference between the light measured in the first slot and the measurement of the second slot as a function of the light received in the first slot. In the figure we see that the difference is proportional to the first measured value and can be approximated by a linear function. This function is then used to correct the second measured value to provide stable values for the same amount of received light. This correction enables precise light measurements and therefore a fast and reliable synchronization.

B. Throughput versus Offered Load

The throughput performance is shown in Fig. 7. The figure presents the results for different values of offered load and distances between sender and receiver. The experiment is performed in the presence of indoor ambient light. Transmitting

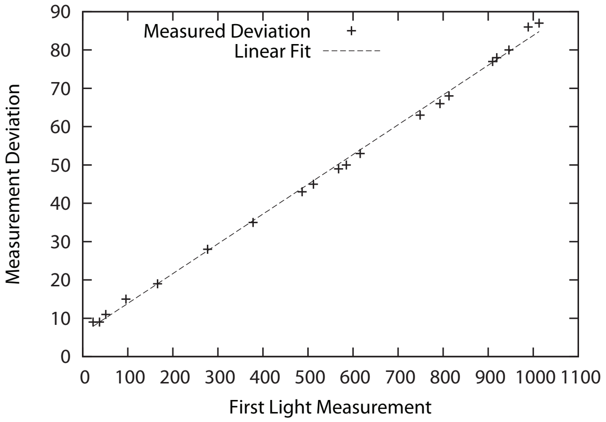

Fig. 6. Deviation of the second measurement of two consecutive measurement with the same amount of light dependent on the light intensity of the first measurement. The measurement are fitted to a linear function that is applied in the implementation to correct the measurement. The computed parameters for the function $f(x) = m * x + b$ are: $m = 0.0777 \pm 0.0009$ and $b = 6.10 \pm 0.5684$ The units shown are conversion values from the ADC, which maps 0 to 5V to a scale of 0 to 1023.

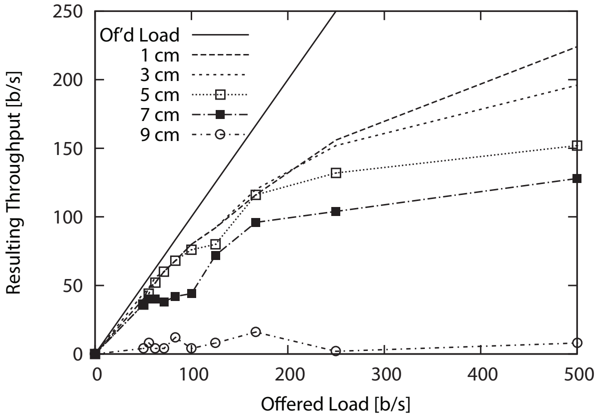

Fig. 7. Throughput versus offered load using the combined light emissions and light measurement approach in a standard office environment for different distances.

and receiving LEDs are placed in front of each other to reduce the effect due to the limited field of view typical of LEDs. The offered load is modified by changing the length of the idle pattern that separates two consecutive messages. The graph therefore shows the impact of the available time to synchronize on the correctly received amount of data. Each data frame is validated using the 2 byte CRC, and only frames received completely correct are accounted for in the throughput. The measurements show that within 5 cm and up to an offered load of 250 b/s, the packet loss ratio is constant. For a distance of 3 to 7 cm we observe that the resulting throughput levels in respect to the offered load for loads heavier than 250 b/s. This indicates that due to the shorter synchronization time and higher distance, more bits are flipped. Nevertheless, the overall linearity shows that the synchronization is fast and more synchronization time has not a big impact. The very low and random results for 9 cm show that the sender and receiver could not synchronize and the received bytes could only be decoded because the two signals just happened to be in phase for some time.

V. Related Work

A survey about VLC and free space optics can be found in [7] and [8]. The usage of VLC for consumer products and toys was shown in [1], [2], and [9]. Some recent standardization efforts are worth mentioning. The Visible Light Communications Consortium (VLCC) aims to publicize and standardize the VLC technology [10]. The IEEE 802.15.7 standard for VLC covers both medium-access control (MAC) and the physical layer (PHY) air interface [5] . The MAC enables peer-to-peer and star topologies, with support of broadcast messages. A slotted random medium access protocol similar to the MAC protocol employed in our prototype is proposed. The 802.15.7 physical layer (PHY) is divided into three types with different modulation schemes and resulting throughput of more than 10 kb/s. We obtain lower rates than this standard as we target a simple implementation with a single LED as receiver. Finally, another VLC standardization effort is driven by [11], which aims towards high-speed communication.

VI. Conclusion

The LED-to-LED communication enables new mobile applications for entertaining and creative use cases and provides a low-cost, low-complexity path to light communication. Visible light is already often part of our environments, so VLC provides a unique opportunity to provide communication capabilities that is not noticed. However, VLC requires cost-effective platforms, and this paper describes a simple design that leverages modern microcontrollers to overcome the limitations of simple, mass-produced LEDs when used for VLC.

References

[1] G. Corbellini, S. Schmid, S. Mangold, T. R. Gross, and A. Mkrtchyan, “LED-to LED Visible Light Communication for Mobile Applications,” in *Demo at ACM SIGGRAPH MOBILE 2012*, 2012.

[2] N. O. Tippenhauer, D. Giustiniano, and S. Mangold, “Toys communicating with LEDs: Enabling Toy Cars Interaction,” in *Demo at Consumer Communications and Networking Conference, CCNC, IEEE*, 2012.

[3] P. Dietz, W. Yerazunis, and D. Leigh, “Very low-cost sensing and communication using bidirectional leds,” in *UbiComp 2003: Ubiquitous Computing*. Springer, 2003, pp. 175–191.

[4] J. G. Proakis, *Digital communications*. McGraw-Hill Series in Electrical and Computer Engineering, 2001.

[5] 802.15.7, “IEEE standard for local and metropolitan area networks. Part 15.7: Short-Range Wireless Optical Communication using Visible Light,” 2011.

[6] Atmel, “8-bit Atmel Microcontroller with 4/8/16/32K Bytes In-System Programmable Flash,” http://www.atmel.com/Images/8271S.pdf.

[7] H. Elgala, R. Mesleh, and H. Haas, “Indoor optical wireless communication: Potential and state-of-the-art,” *IEEE Communications Magazine, Optical Communications Series*, vol. 49, no. 9, pp. 56–62, 2011.

[8] ——, “Indoor broadcasting via white LEDs and OFDM,” *Consumer Electronics, IEEE Transactions on*, vol. 55, no. 3, pp. 1127 –1134, 2009.

[9] S. Schmid, M. Gorlatova, D. Giustiniano, V. Vukadinovic, and S. Mangold, “Networking Smart Toys with Wireless ToyBridge and ToyTalk,” in *Poster Session, Infocom*, 2011.

[10] VLCC, “Visible Light Communications Consortium,” http://www.vlcc.net, viewed 12-01-2012.

[11] Omega, “the Home Gigabit Access project,” http://www.ict-omega.eu/, viewed 12-01-2012.