

Frank Englert, Patrick Lieser, Marius Rettberg-Päprow, Alaa Alhamoud, Doreen Böhnstedt, Ralf Steinmetz:
LightMon: Apportioning the Effect of Light Switching Events on the Electricity Consumption of Buildings. In:
EWSN '16 Proceedings of the 2016 International Conference on Embedded Wireless Systems and Networks,
February 2016

Frank Englert, Patrick Lieser, Marius Rettberg-Päprow,
Alaa Alhamoud, Doreen Böhnstedt, Ralf Steinmetz
Multimedia Communications Lab
TU Darmstadt
Darmstadt, Germany
{firstname.lastname}@kom.tu-darmstadt.de

Abstract

A major share of the electricity consumption in buildings is caused by the lighting installation. In contrast to other consumers of electricity, the components for lighting are distributed over the entire building. Therefore, measuring the electric energy spent for lighting with the appropriate level of detail is a very costly endeavor. In order to mitigate the underlying problem, we developed the LightMon system which correlates light switching events obtained from a Building Control Network (BCN) with power measurements on the mains distribution level. Combining both sources of information allows our system to estimate the electricity consumption for each light individually with an estimation error below 11.34%. The system adapts automatically by observing switching events and their effects on the power consumption. This adaptation occurs automatically within 8 to 14 days during normal building operation - no manual actions are required during this training phase. In addition to the above mentioned characteristics, the system can be deployed in most buildings with commercial use. The LightMon system facilitates comprehensive information gathering on the electricity consumption of buildings, enabling advanced energy optimizations.

Categories and Subject Descriptors

[**Computer systems organization**]: Embedded and cyber-physical systems—*Sensors and actuators*

General Terms

MEASUREMENT

Keywords

Building Control Network, Load Disaggregation, Electricity Consumption, Buildings, Lighting

1 Introduction

In the United States (US) as well as in the European Union (EU), approximately 11% of the total electricity produced is spent on lighting [7, 9]. In buildings, lighting accounts for 30% to 49% of total electricity consumption [28]. This results in significant financial and environmental costs for our society. However, due to high setup cost, energy efficiency technologies are often not used although significant savings could be achieved. Recently developed technologies such as Solid State Lighting have the potential to reduce the electricity demand of lighting by a factor of four. In combination with adaptive lighting systems capable of sensing the occupancy as well as the current luminance conditions (c.f. [25, 26]) it becomes possible to achieve optimal lighting conditions at the lowest possible energy expenditure. However, the adoption rate for these methods remains low. The key obstacles these technologies face are firstly the high investments required and secondly the general lack of knowledge regarding to the vast saving potentials. Current Building Control Networks (BCN) like KNX, Dali, EnOcean do not mitigate this blind spot as those systems typically do not monitor switched loads. The integration of dedicated electricity meters would increase the cost of those systems and is therefore often omitted.

LightMon was developed with the consumer-centric aim of providing deep insights by making detailed information on the costs of electric lighting easily available. Our system obtains light switching events from the building control network (BCN) and correlates them with changes in the whole building's power consumption. Our LightMon system is easily deployable in buildings which are already equipped with a BCN system for controlling the lights. This requirement is fulfilled for most new commercial buildings as BCN simplifies electric wiring and increases the flexibility. By comparing the whole building power consumption before and after a switching event, LightMon creates a comprehensive electricity consumption database stating the power draw for each light individually. In conjunction with historic switching events, it becomes possible to calculate the electricity cost for each light individually, on room level granularity or for the whole building. During a one year field study, LightMon estimated the electricity consumption of 69 switched lamps with an error below 11.34%. LightMon is based on adaptive and scalable algorithms, which achieve live processing even on low-power embedded systems.

Main Challenges

Detecting the correlation between switching events and power changes with an appropriate level of detail faces the following main challenges:

1. Cost Pressure: Due to high cost pressure, nodes of BCN are kept as simple as possible. Therefore, BCN switching actuators typically do not contain additional hardware components for current measurement of the attached loads. In order to keep a low resource profile, common smart meters have limited sampling rates below 1 Hz [4]. Electricity meters with higher sampling rates are often specifically tailored for research purposes [23, 14, 21] and not widely available in today’s buildings. Electricity meters providing higher sampling rate would require more processing power as well as memory and thus manufacturers omit this functionality.

2. Concurrent Events: High probability of concurrent events and turn-on-transient convolve the effects of individual events. Thus, differentiation is not sufficient to handle conflicting situations. A high number of collisions arise from the fact that buildings typically share a BCN for all sensors and actuators. While manual switches are rarely used, passive infra-red sensors (PIR) trigger switching commands on every detected movement which causes high number of events. In the building hosting our observed office environment with 5 floors, 10 hallways, and underground parking there is a probability of 43% for colliding events within a two second time interval, as shown in Figure 1. This Figure shows the occurrence probability of different inter-event-times as well as the cumulated density function (CDF). The underlying excerpt of data was collected over one month during our field study and includes weekends as well as night times. This figure implies to use electricity meters supporting high frequency sampling. However, this approach is not sufficient as the inrush transient of lamps from cold start is significant (c.f. Sec. 4.4). Therefore collisions could not only be mitigated by using electricity meters with high sampling frequency.

3. Noise Level: Other electrical appliances which also consume electricity obfuscate the effects of switching events on the whole building power consumption. Most electrical appliances have variable power consumption based on their current mode of operation. Typically, those appliances are not connected to the BCN and thus their mode of operation remains unknown. The variability in the power consumption caused by those devices interfere with effects of switching events and thus must be handled by LightMon.

Our proposed system copes with those challenges in the most cost-efficient way. Combining low sampling rate requirements on the electricity metering side with plug-in integration in existing BCN systems, our proposed solution provides cheap means to obtain key insights in the electricity spent for electric lighting.

Contributions

In order to address those main challenges, our work provides the following contributions:

1. We developed simple, adaptive and scalable mechanisms for estimating the effect of every single BCN light switching event on the total electricity consump-

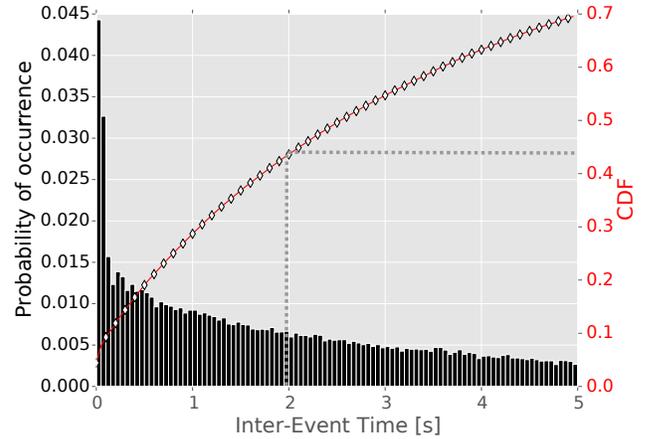


Figure 1: Event distribution observed in office building.

tion. Our system handles concurrent switching events and achieves live-processing capabilities on resource-constrained hardware.

2. LightMon is easily installable in modern office buildings by deploying a single-box system which interfaces the BCN and measures the power consumption on mains distribution level.
3. We provide a thorough evaluation based on a simulation as well as a real-world field study in our institute with 23 offices over the duration of one year.
4. We make the entirety of the data collected during our evaluation publicly available. This dataset collected during our field study contains the electricity meter readings of two electricity meters, ground truth data for 69 installed lights as well as all BCN switching events which happened in the building.

Benefits

In the following section we present four use-cases benefitting from LightMon:

1. Cost visibility: Our system exposes electricity cost of lighting for single lights, on a room level, or for the whole building. This way, it becomes possible to identify hot spot configurations with excessively high electricity consumption and thus high costs.

2. Lamp selection: Additionally, our system provides insights into historic usage patterns like usage frequency, the average switch-on durations, or total daily usage, which are vital parameters for selecting the most energy- and cost-efficient lighting technology.

3. Fault detection: Our system enables the automatic detection and localization of faults in the lighting system, such as broken lamps. This reduces the response time to faults and enables more efficient scheduling of maintenance operations by facility managers.

4. Adaptation: LightMon is an important information source for systems which aim for adaptive behavior. Our system enables quantifying the saving potentials of strategies which 1) enhance the energy efficiency of lights, 2) reduce the supply of light to spaces where no-one is present, and 3)

limit the over-provisioning of lighting. Furthermore, the information is useful for implementing demand response (DR). The system could slightly adapt dimmers when receiving load shedding requests. If the luminance change is below a certain threshold, the variation would not even be noticed by inhabitants of the environment. If, for example, nation-wide retailers with hundreds of stores adapt such a DR schema, LightMon could be used to easily assess the DR potential.

The remainder of this work is structured as follows: First, we describe the most important related works in Section 2. Building upon that, a concise description of our LightMon system is provided in Section 3 alongside a thorough evaluation in Section 4. Next, we lay out the insights gained from the application of LightMon in our very institute’s building. Finally, we supply a conclusion to this paper.

2 Related Works

Electricity metering on a device-level granularity is an important foundation for optimizing energy efficiency in buildings. Having this information allows the collection of detailed feedback on the amount and originating appliance of all electricity consumption. According to Darby [10] and Fischer [15], this kind of user feedback alone causes behavioural changes which lead to significant electricity savings. These energy reports can be enriched even further with links to contextual information. E.g. by using apportionment schematics [17, 27], the consumed electricity can be matched with individual inhabitants. Fused with contextual information like user activities [2] or locations of the inhabitants, home electricity saving recommender systems can be implemented [1].

However, electricity metering on a device-level granularity is challenging. Beginning with the works of Hart, the research community has explored numerous approaches for this kind of sensing task. Related works in this area can be categorized into two main directions: those propagating (1) the deployment of additional sensors or (2) using load disaggregation algorithms. When deploying additional sensors, each monitored appliance needs its own sensor unit. In such a setup, one can either directly measure the current and voltage of the attached appliance (ACme [19], DeltaFlow [8]). Alternatively, other indirect effects like electromagnetic emissions [16], audio signals [12], inhabitants position [29] or heat can be observed to infer the electricity consumption of the device under observation (Kim et al. [22]). For some devices it may be possible to embed energy models directly to the device firmware for inferring the current power draw for the current mode of operation [13]. The deployment efforts of such systems scale linearly with the number of appliances to observe. In contrast, load disaggregation approaches aim for a much lower deployment effort [23]. The so-called Non-Intrusive Appliance Load Monitoring (NILM) relies on a single electricity meter installed at the mains distribution, which measures the summarized power of all connected appliances. The resulting power signal is split into its components using a disaggregation algorithm. As described by Beckel [5], the design space to implement a NILM algorithm has three dimensions: (1) power model representation, (2) device state detection, and (3) model training. Differences in the exact behavior of present appliances, the amount of power

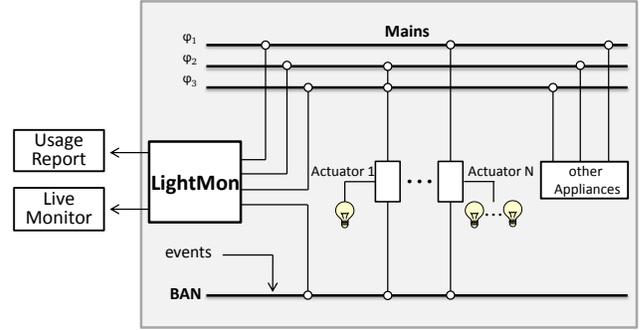


Figure 2: Integration of our System into an existing building infrastructure.

consumption and the number of observed appliances, require different trade-offs. To compare different load disaggregation algorithms, the decent NILMTK [4] benchmarking toolkit can be used. However, using pure NILM approaches, the disaggregation of low loads or environments with many similar appliances remains challenging [3].

Therefore, hybrid approaches combine NILM algorithms with information obtained from easily deployable environmental sensors in order to enhance the disaggregation accuracy (c.f. Pathak et. al [24] with audio as information source, Irwin et. al [18] integrating home automation systems, or [20, 6] with general on/off events). This information significantly reduces the uncertainty on the set of present appliances or their internal working state and therefore enhances the accuracy of load disaggregation.

3 Calculating the Energy Demand of Switching Events

The goal of this work is to describe the means for determining the electricity consumption of the electric lighting in office buildings with a high level of detail. Our two-phase approach relies on the existence of a building control network (BCN) which is used for switching the lights. As those BCNs are commonly used in modern office buildings, this requirement is barely a limitation.

The working principle of our system called LightMon is shown in Figure 2. To determine the electricity demand of the lights, our system continuously monitors the electricity consumption of the whole building using electricity meters installed at sub-distribution boxes. Typically, each building has one or more sub-distribution boxes per floor. Additionally, our system records all light-switching events that are exchanged over the BCN. During an unsupervised training phase, our system correlates switching events observed from the BCN with the actual electricity consumption obtained from the electricity meters. This way, the influence of single switching events on the data collected by the electricity meters can be inferred. During the training phase, all switched actuators are annotated with their inferred influence on the various electricity meters installed in the building. As a result, we obtain a power consumption matrix stating the individual power draw of each lamp that was switched at least once. We assume that each light in a building is attached to a BCN actuator and that each actuator has a unique address. More

precisely, the BCN switching event is sent to the actuator, which then connects or disconnects the light from mains. It is common practice to connect multiple light bulbs to a single actuator or to switch multiple actuators using a group address. Therefore even though light bulbs with the same power rating are used, the load varies significantly for different switched actuators as a different number of lights is switched. For better comprehensibility, we define the term "switched light" as a BCN actuator to which one or more light bulbs are connected.

In the reconstruction phase, our system uses this power consumption matrix in conjunction with a stream of observed BCN light switching events to reconstruct the share of the power consumption caused by the lighting. However, this approach is not limited to whole-building observations. By selecting the switching events of single rooms or even single switches, our system calculates the electricity demand for lighting in much greater detail. For example, we can use our system to calculate the lighting electricity demand for each of the 23 offices in our department building independently.

Having given a brief overview over of working principles of our system, we shall continue by describing the underlying algorithmic steps in the following paragraphs.

3.1 Definitions and Assumptions

The implementation of our LightMon system is based on seven core assumptions. These are explained below:

1. We assume the availability of entire building power measurements with a sufficiently high sample rate.
2. If a building is supplied using a three-phase power supply, we assume the availability of independent power measurements for each available phase. This assumption holds for most commercial buildings. In our observed office environment, independent measurement points for six phases were available.
3. We assume the usage of a building automation network to switch lights. Supported switch commands are on, off, or any level in between (dimmer functionality). Newer commercially used buildings are equipped with BCN systems like KNX, EIB or Dali.
4. As our algorithms are executed on embedded system hardware, we assume a limited availability of computation- as well as memory resources. We trade accuracy in favor of simpler algorithms, which still achieve good results and even run as "live systems".
5. We use the heavyweight function as a power model for our light bulbs. Thus, we assume zero power draw in "switch off" mode and a constant power draw if a particular light is switched on. This model fits well for LEDs or incandescent lamps. However, CCFL lamps need a certain warm-up time to reach steady state.
6. If a particular lamp is not operated at maximum brightness level, we assume a linear correlation between current brightness level and power consumption (c.f. Roisin et. al [26]).
7. The residual noise $R(t)$ measured on the power signal is bounded below a certain threshold (c.f. Section 4.3 for estimates of this threshold).

Table 1: Symbols which are used for modelling the working principle.

Symbol	Description
Event e	3-tuple of (sender, actuator, cmd)
Power $p_{e,i}$	Influence of e for the actuator i on the total power p in Watt on the n 'th phase φ_n
N	Number of independent mains lines in power supply system of the building
I	Total number of switched actuators
Phase φ_n	n 'th independent mains line for n in $1 \dots N$
Time t_e	Timestamp t at which event e occurs
Residual $R(t)$	Power draw of all, non-observable loads present in the building
State $x_i(t)$	Switch state of actuator i at time t
Vector Φ_e	Phase Effect Vector of event e in Watt $(p_{e,1}, p_{e,2}, \dots, p_{e,N})^T$ for phase $1 \dots N$
Vector ε	Measurement error vector in Watt $(\varepsilon_1, \varepsilon_2 \dots \varepsilon_N)^T$ for phase $1 \dots N$

In addition to these assumptions, Table 1 explains important symbols required to understand the exact working principle of our event correlation algorithm.

3.2 Training Phase

In the training phase, we derive the power consumption matrix from effects of switching events on the power consumption of the building. Depending on the intended use cases, the training phase could run continuously or only for predefined time windows. When running continuously, it becomes possible to detect faults or anomalies in the lighting system of the building. When scheduled for pre-determined time windows, it becomes possible to further reduce the computation requirements of our system. As shown in Figure 3, the training phase consists of five processing steps, which are described in the following paragraphs.

1. Delay compensation: In order to correlate switching events with electricity consumption, both data streams must be synchronized. Due to filtering in the electricity meter, the measurements of instantaneous electricity consumption lag behind the occurrence of BCN events by a few seconds. This static delay depends on the exact hardware configuration and must be compensated for. To determine the delay, we extract switching events from the electricity time series (c.f. [11]). Next, we calculate the average time difference between switching events observed from the BCN and the corresponding switching events extracted from the electricity time series. Using our hardware configuration, we observed a maximal time delay $t_{delay} < 2sec$.

2a. Naive handling of conflicting events: If multiple BCN switching events happen within a small time window, their respective power consumptions interfere with one another. Thus, the observed effects on the building's electricity

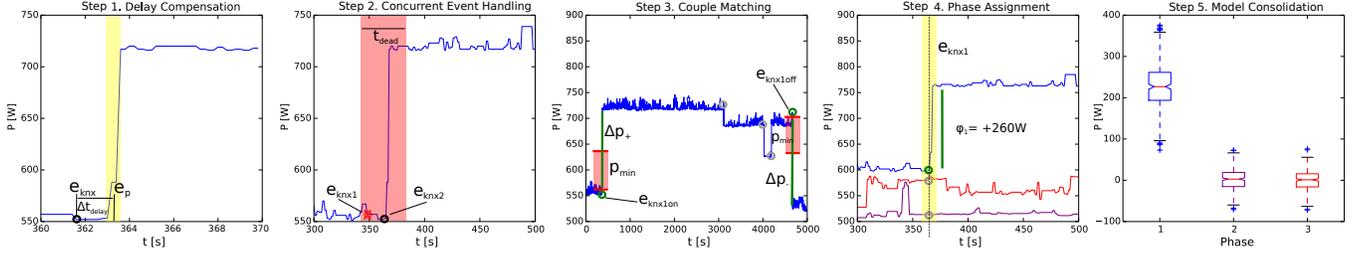


Figure 3: Process steps to calculate the influence of electric switching events on the total electricity demand.

consumption cannot directly be assigned to the exact event which caused the change. We implemented two methods for handling such situations. The naive approach relies on dropping conflicting events whereas the second method uses a linear equation system to gain information even from conflicting events. For the naive approach, we define a time window t_{dead} around each BCN switching event in which no other event may occur. If multiple events happen within the time window t_{dead} , all those events are excluded from the training phase, as their influence cannot be determined unambiguously. However, this filtering step comes at a cost. If BCN switching events have a temporal correlation, their influence on the electricity consumption can never be calculated if filtering occurs. This occurs if several hallway movement sensors trigger the same light on movement. Furthermore, the maximum number of BCN switching events which can be processed is limited by t_{dead} . Especially in large environments with hundreds or even thousands of actuators and a high number of switching events, the likelihood of collisions becomes very high. Thus, to achieve good scalability, the parameter t_{dead} should be as small as feasible. Yet, t_{dead} cannot get arbitrarily small. Its minimum size is determined by two factors. First, it depends on the gradient of electricity consumption change in case of switching events. Second, if the electricity meter has a low sampling rate, larger values for t_{dead} are necessary to acquire a data points before and after the switching event. Experimental methods to determine optimal values for t_{dead} are presented in Section 4.

To further minimize the number of excluded events, our system uses the current power consumption matrix to check if multiple events occurring in the same time window t_{dead} would influence the same electricity metering unit. If the switching events are known to influence different metering units and thus different parts of the building, these concurrent events do not cause statistical uncertainty. Of course, this optimization only works if the current power consumption matrix is already populated.

Using the naive conflict handling strategy, the effects of a BCN switching event on the electricity consumption is calculated as follows: As assume the switched light reacts to switching commands with a step function, the influence $p_{e,i}$ of event e and actuator i is:

$$p_{e,i} = p(t_e + \Delta t) - p(t_e - \Delta t); \Delta t = 0.5 t_{dead}$$

In other words, we subtract the power consumption before the switching event e from the power consumption after the switching event. The time span Δt directly depends on t_{dead} . In case of switch-off events, we invert p_e in order to obtain

positive values for p_e . In time frames with high variations of electricity consumption, p_e may become negative. These values are indicators for invalid measurements and thus filtered. The information content of on- and off-events is not equal. The inrush of some electricity-saving lights increases slowly during a warm-up time. Thus, switch-off events are better for approximating the steady state power draw of these lights. However, as the switch-on behavior is deterministic, we calculate the ratio between on- and off-events and use this information to scale up the on-events (c.f. assumption 5). While it may be possible to use more sophisticated methods for modeling the inrush, we expect the lamps to run in a steady state for most of the time.

2b. Enhanced conflict handling using superposition:

A better approach towards conflict handling is imposed by the particular problem structure. The general idea is using a linear equation system to describe the superposition of power demand in case of concurrent events. In detail, the system works as follows: The total power consumed by the building is equal to the summarized power draw for lighting plus the power draw of all other consumers which are switched on:

$$P_{tot}(t) = \sum_{i=1}^I p_{e,i} * x_i(t) + R(t)$$

By differentiating this term with respect to time t we get the change term:

$$\frac{d}{dt} P_{tot}(t) = \frac{d}{dt} (p_{e,1} * x_1(t) + \dots + p_{e,I} * x_I(t) + R(t))$$

If only a single state change happens in each time window, this term is equal to the naive conflict handling. In this case, all elements without state change drop out as their derivative is zero. Then the respective value of p_e is equal to the power change ΔP . However, in case of multiple concurrent state changes their effects on the power consumption superimpose. For decomposition of the power signal, we use the following linear equation system:

$$\frac{d}{dt} \begin{pmatrix} x_1 & x_2 & \dots & x_I & t_{e0} \\ x_1 & x_2 & \dots & x_I & t_{e-1} \\ \vdots & \vdots & \ddots & \vdots & \vdots \\ x_1 & x_2 & \dots & x_I & t_{e-I} \end{pmatrix} \begin{pmatrix} p_{e1} \\ p_{e2} \\ \vdots \\ p_{eI} \end{pmatrix} = \frac{d}{dt} P(t) - \frac{d}{dt} R(t)$$

To solve this equation system, at least I linearly independent state changes are needed. For populating the equation system, the last I linearly independent events are taken. As result we obtain the power influence vector p for each actuator. In a

multi-phase system, this step must be repeated for each phase individually in order to obtain the phase influence vector Φ_e for all actuators. If $R(t)$ is constant, there is an exact solution of this equation system. Otherwise, the presence of variations limits the performance of our system. In section 4.3 we quantify the impact of varying $R(t)$ on the overall system performance.

3. Couple Matching: In addition to the conflict handling strategy, we apply a matching of corresponding on- and off-events for each light switched by the BCN. Using a simple state-machine, we filter duplicate switching commands that cause no changes to the environment and thus have no effect on the electricity consumption. However, purely event-driven BCNs like KNX provide no functionality to query the current state of a light. Thus, when our system starts processing events, it cannot determine the current state of lights present in the building. To cope with that fact, our state machines are modeled with switched-off as their initial state. Initially, this may cause false negatives when the light's real state diverges from our state machine. After receiving the first switch command, both states are in sync, however.

4. Phase Assignment: In Europe, the electric wiring in a building is connected to the three phase electricity grid. Therefore, our proposed system measures the electricity consumption for each phase independently. During our field study in an office building, our two electricity meters observed the power draw of six phases independently. To achieve a similar load on all phases, different parts of the building are connected to different phases. Therefore, BCN switching events typically do not influence all phases in an environment. In this processing step, our system determines which phases are influenced by BCN switching events. To estimate the Phase Effect Vector Φ_e for a switching event e we calculate p_e for each phase and check if p_e is above a certain threshold p_{min} :

$$\Phi_e = \begin{cases} p_{e,n} & \text{if } p_{e,n} > p_{min} \\ 0 & \text{else} \end{cases} \quad \text{for phase } n \text{ in } 0 \dots N$$

For each observed environment with N phases, Φ_e is an N -dimensional vector describing the influence of the event e on the electricity consumption of each phase in the system. As some BCN switching events do not cause effects, p_{min} suppresses the influence of random variations. Random variations are caused by other electrical appliances which are also connected to the power meter. The threshold p_{min} determines the system performance. If the value for p_{min} is too small, a power consumption p_e may get assigned to events without effects on the electricity consumption. On the other hand, if p_{min} is too large, the effects of events caused by switching actuators with a power draw below p_{min} are not detectable by our system. While a good trade-off for p_{min} depends on the exact environment, we show means to approximate p_{min} in Section 4.

5. Model Consolidation: The goal of this processing step is to enhance the accuracy of the power estimate for a switching event e by combining multiple observations of its effects Φ_e . Each single observation consists of the real effect of e $\Phi_{e,real}$ and a measurement error vector ϵ :

$$\Phi_e = \Phi_{e,real} + \epsilon$$

By consolidating multiple observations of Φ_e it becomes possible to prune instances with a high error ϵ and also to estimate $\Phi_{e,real}$. In order to remove erroneous instances, we use a sliding-window based approach to merge the last W observations of Φ_e . The window containing the last W observations of Φ_e is called Consolidation Matrix:

$$CM = \begin{matrix} & \varphi_1 & \varphi_2 & \dots & \varphi_N \\ \Phi_{e1} & (p_{e1,1} & p_{e1,2} & \dots & p_{e1,N}) \\ \Phi_{e2} & (p_{e2,1} & p_{e2,2} & \dots & p_{e2,N}) \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ \Phi_{eW} & (p_{eW,1} & p_{eW,2} & \dots & p_{eW,N}) \end{matrix}$$

Using this consolidation matrix, we consecutively apply three different strategies (1) Blanking, (2) Clustering, and (3) Averaging to reduce the effect of erroneous instances and finally to approximate $\Phi_{e,real}$. The blanking strategy enhances the robustness against random noise on phases which are unaffected by a switching event e . For each column φ_n in our consolidation matrix, the number of zero entries is determined. If a certain percentage of all entries are zero, we assume no effect of event e on the phase φ_n , and thus, the whole column is set to zero. However, it is still possible that entries in the consolidation matrix for event e_w significantly deviate from the real power draw of a switched actuator. The reasons for this may be random variations happening on the phase φ in the same time as event e_w . It is unlikely that random variations for different events have the same magnitude. In contrast, events without significant noise components ϵ form a cluster around $\Phi_{e,real}$. To identify those entries with high noise variations, we apply a DBScan clustering with a sufficiently low bandwidth parameter on consolidation matrix CM . We assume the biggest cluster to be found around the real power value $\Phi_{e,real}$. Therefore, we set all values to zero which are not assigned to the biggest cluster. Finally, we consolidate our matrix by summing up the average power values for each phase φ :

$$\Phi_{e,est} = \text{mean}(CM, \text{column wise})$$

Using a sliding-window based approach to consolidate the last W events has two major advantages. The first advantage is a bounded computation complexity. As the window size W is known in advance, it becomes possible to ensure that our algorithms do not exceed the available computation or memory resources on a embedded hardware platform. As a second advantage, it becomes possible for our system to adapt to changes by forgetting old values. If the power consumption of a light changes, (e.g. because one of many light bulbs in a switching group breaks), this new behavior will be reflected in the power consumption matrix after a few events.

4 Evaluation

To assess the applicability of our system we conducted an exhaustive evaluation of our system. For rating the quality of our system we focus on the performance criteria:

1. Accuracy
2. Scalability
3. Robustness against noise
4. Low Resource Consumption

Our evaluation consists of a simulation part to measure performance criteria and a field study in an office building to verify the simulated results. As it is not feasible to equip all lights in a building with power meters we use a simulation to create synthetic environments with different properties. More precisely we use historic power traces obtained from the sub-distribution of an office building and added virtual BCN switching events. The historic power trace was trimmed to a time window from 7a.m. to 6p.m. during a normal working day. This way we ensure realistic noise conditions on the power signal for our simulation. Night times, weekends and holidays were excluded from our trace set, as these days contain too few variations in the power consumption. To simulate the effects of the virtual switching commands, we add constant power consumption as long as the simulated light is switched on. Using this approach we can freely vary relevant environmental parameters like the power draw of switched lights (1), the number of different lights present in the environment (2), the total number of BCN switching events, (3) and the noise level caused by other appliances (4). A complete list of all parameters available for simulation is denoted in Table 2. Additionally, the simulations allow searching optimal values for the system parameters p_{min} , t_{dead} or the consolidation strategy. However, to testify the correctness of this aforementioned simulation we conducted measurements in an office environment over the period of one year. This field study is described in section 4.6.

4.1 Accuracy

The accuracy of our system depends on two factors. First, in each environment there are different lights in different switching groups. Our system should be capable of estimating the power consumption independently from the actual configuration. More precisely the error of the estimation should not depend on the power consumption of the switched actuator. Second, as some lights are switched few times a day, the number of events for training the power model should be as low as possible.

To determine the effect of the actuator’s power consumption on the estimation error, we simulated actuators with different power consumptions by varying the parameter act_power from 40W to 1,000W. The resulting estimation error is shown in Figure 5. The graph clearly shows that the total estimation error is nearly constant. Even small consumers with a power draw of 50W could be estimated with a relative error below 4%.

Additionally we tested the accuracy of the resulting power model for a varying number of switching events. This way we measure how the accuracy of our system evolves over time when more and more switching events per actuator are available. To carry out this simulation, we varied the parameter num_events from 2 to 30 and observed the mean absolute estimation error. However, it is important to mention that our simulation environment always generates on/off event pairs in order to avoid situation where sequences of multiple on- or off-events occur. The estimation error for a different number of training events per switched appliance is shown in Figure 6. As visible in the graph, the error drops fast with a rising number of event pairs. At the turning point of ten event-pairs error flattens and more training provides only small benefits.

Table 2: Parameters used for simulations.

Parameter	Value	Description
historic_trace		Historic power recording to use as source.
start_time	7a.m.	Start time of simulation.
stop_time	8p.m.	Stop time of simulation.
num_repeats	30	Number of times the simulation was repeated.
num_acts	50	Number of different actuators present in the environment.
num_events	10	Number of events to execute for each switching actuator.
act_power	50W	Power consumption of the actuator in the simulation.
noise_level	3W	Level of noise to add to the real world power trace.
p_min	30W	Effects of events below this threshold are not considered.
bcn_delay	3s	Time lag until BCN events arrive.
delta_t	2s	Time to look ahead for effect calculation.
b2e_delay	3s	Time lag between BCN events and power readings.
t_dead	4s	Dead time around events.
blank_th	0.3	Percentage of zero p_e required to blank effects of event e on phase ϕ .
a_phase	1	Amplification factor multiplied with the power readings.
W	10	Size of the sliding window to consolidate the last W events for a particular light.

Thus, ten observed event-pairs per actuator are sufficient to derive a high accuracy power model. We selected the size of our sliding window W based on those findings. As side note it is interesting to mention the reason for the drop in the estimation error in case of four training events. Beginning with this number of events the consolidation step starts to remove errors caused by noise on phases not affected by an event e .

4.2 Scalability

The *Scalability* quantifies the total amount of events our system can compute. In general, the scalability of our system is limited by noise effects on the phase. With disabled collision handling, the *Scalability* is also limited by the dead time between two events affecting a single phase. If multiple events happen within t_{dead} , our system removes all those events because it is impossible to determine the effect of each event independently. With a rising event density, the number of conflicting events increases. If a massive number of conflicts occur, the number of training events per actuator

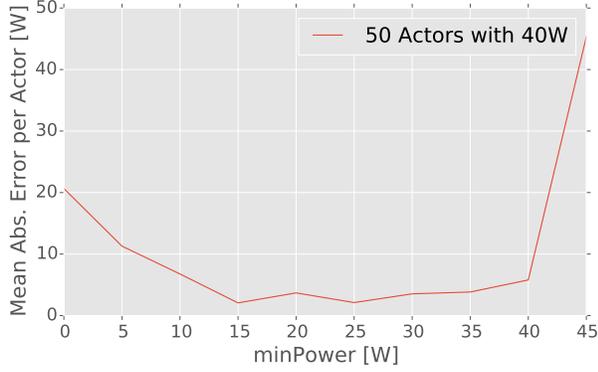


Figure 4: Resulting estimations errors depending on the value of p_{min} .

decreases and thus the estimation error rises. These effects are mitigated using our proposed collision handling schema. In case of N BCN events occurring at the same time, those events are combined with N other, linear independent events to assess the phase influence of each individual event. However, the information density of each event is still reduced by a factor of N .

In order to determine the *Scalability* of our collision handling schema, we varied the number of different actuators and measured the resulting absolute error per actor. As baseline measurement, we drop all events in case of collisions. The results of this simulation are shown in Figure 7. Below 300 actors present in the environment, the number of collisions is rather low and thus both approaches perform equally well. However, with a rising number of present actuators and thus a higher collision probability, the system performance degrades rapidly without collision handling. On the other hand, if collision handling is enabled, the system performance degrades much slower with a constant pace.

To assess the *Scalability* of our system, we measure the error caused by conflicting events for different combinations of switched actuators and events per actuator during a simulated time frame of 12 hours. The result of this simulation is shown in Figure 8. The x-axis shows the number of different actuators used for the simulation whereas the y-axis denotes the number of events per actuator. For each tuple (x,y) we calculated the resulting error. Low errors are indicated by green color (bottom and left area of Figure 8) and areas with high error are marked red (c.f. top right corner in Figure 8). As expected, the estimation error increases with a rising number of events. However, below an event density of 400 events per hour, the resulting error is below 5W, which is considered as low. For a high number of different actuators and a low number of events per actuator the error increases due to the insufficient number of events for training.

4.3 Robustness against noise

The total error ϵ_{total} for inferring the power consumption of event e is dominated by two main components. Namely,

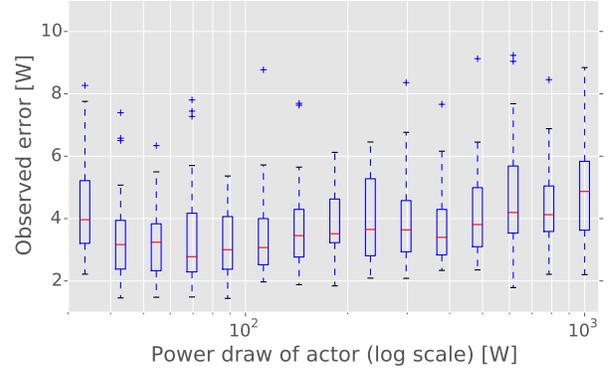


Figure 5: Estimations are independent from the electricity demand of actuators.

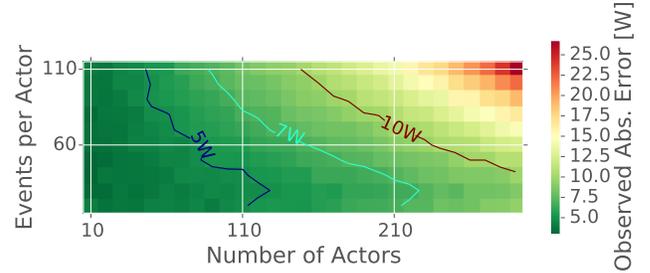


Figure 8: Influence of event count on estimation error.

these are event estimation error ϵ_{event} and phase error ϵ_{phase} :

$$\begin{aligned} \epsilon_{total} &= \epsilon_{event} + \epsilon_{phase} \\ \epsilon_{event} &= \epsilon(\phi_n) \\ \epsilon_{phase} &= \sum_{i=0, i \neq n}^N \epsilon(\phi_i); \quad \text{if } \epsilon(\phi_i) > p_{min} \end{aligned}$$

The error ϵ_{event} is caused by inaccurate estimation of the effect of event e . The component ϵ_{phase} is caused by random variations above p_{min} on other phases which are co-occurring with event e . For noise below p_{min} the phase error is 0, as it is completely removed by the blanking strategy of the consolidation step (c.f. Section 3.2). With increasing noise level on the observed phases, the phase error starts to dominate the error term, as the blanking strategy could no longer successfully identify and prune erroneous instances. We simulated the effects of noise on the accuracy of our system by adding artificial noise. This noise signal is generated using a logarithmic series distribution with $\alpha = 0.85$. Our findings with real power traces indicate that the medium tailed distribution of the logarithmic time series provide a good approximate to the noise observable in our test environment. In order to test our system with different noise levels we varied the parameter *noise_level* to scale the noise signal with different noise amplitudes from 0W to 30W. The results of this experiment are shown in Figure 9. From 0 to 10W noise level the error ϵ_{total} is dominated by ϵ_{event} and correlates nearly linear with the noise amplitude. Above 10W ϵ_{total} is dominated by ϵ_{phase}



Figure 6: Approximately eight switching events are sufficient for training.

which results in a much higher slope as we simulated an environment with six phases. In our field study environment we observed a mean noise power below 3W and a noise shape similar to a logarithmic distribution. Therefore the result of this experiment shows that our system could deal with a five times higher noise level without raising the relative estimation error of a 100W actuator above 6%.

4.4 Simulation Accuracy

While being conducted with great care, our simulation could not address all variabilities which might be observed in real world deployments. Effects not considered by the simulation are:

1. Turn-on-transient of lights
2. Ageing of lights
3. Jitter on the BCN

Those effects and their impact on the accuracy of LightMon will be discussed in the following section.

Our simulation framework assumes immediate response of lights to switching events. While this is a valid assumption for LEDs, CCFL-lights as well as incandescent light bulbs have different turn-on behaviour. To determine those transients, we measured the inrush power of a cold cathode fluorescent lamp (CCFL) as well as a halogen incandescent light bulb using a high frequency power meter. As illustrated in Figure 10, halogen light bulbs respond within 200 ms to switching commands while CCFL lights take up to 2 s to reach steady state power consumption. As the shape and also the inrush duration varies for different lights, we decided to exclude the turn-on-transient from our simulation. Therefore, if the turn-on-transient of a light exceeds t_{dead} , LightMon could not accurately track the power consumption. Most probably, this fact accounts for discrepancies between simulated values and measurements we obtained in a real world setup.

The inevitable ageing effects reduce the performance of lights over time. However, different lighting technologies have wear-down time frames ranging from 1.000 hours mean-time-to-failure (MTTF) for incandescent light bulbs to 30.000 hours for modern LEDs. However, as the ageing happens slowly over time, LightMon adapts its coefficients within a few switching cycles to accurately track the real power draw of the attached lights. At some point in time a light may finally

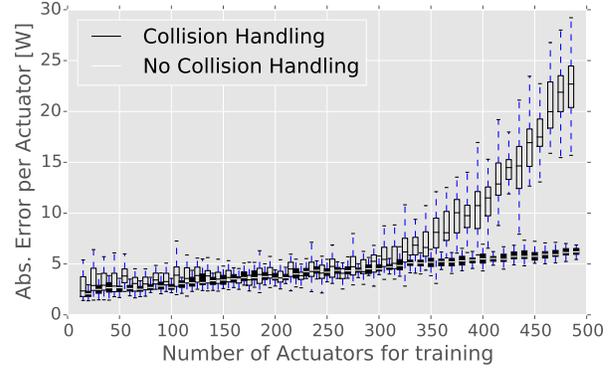


Figure 7: Influence of actors count on system performance with and without collision handling

fail, which could be detected and reported by LightMon.

Additionally, in high load situations on the BCN, there may be a jitter effect on the otherwise static time-lag between switching events and power signal. This jitter could cause temporal de-synchronization of BCN events and their corresponding effects on the measured power. This leads to false power assignments for switching events and degrades the system performance. Currently, our simulation does not account timing jitter in the arrival of BCN events. Jitter below a certain threshold is non-critical while high jitter values could be avoided by proper configuration of the BCN.

4.5 Low Resource Consumption

We measured the performance of LightMon on a 700MHz Raspberry Pi B+ by replaying recorded historic electricity consumptions as well as the corresponding BCN events. The replay data covers five working days. The replay log contained a total of 1,453 switching events. Processing of that data with the steps described in Section 3 took between 126 and 130 seconds. In average our system running on a Raspberry Pi B+ is capable of processing 11.3 switching events per second. Our current implementation is running on Python 2.7 without any optimizations for better performance. The Raspberry Pi which executes LightMon consumes between 1.5W under light load and not more than 3W under full load. Together with the combined power draw of 4.8W for our electricity meter, our system consumes below 7.8 Watt. Compared to the electricity demand of our institutes building, the installation of LightMon increased the electricity demand only by 0,5%.

4.6 Field Study

To testify the functionality of our LightMon system, we installed the system in the second floor of a four story office building. The system has been installed there for over one year and monitored 69 light switching groups in 23 offices, one meeting and one resting area. During that time, the office area was in unrestricted use by approximately 45 people. Among lighting equipment, the power meters observed the power consumption of Laptops, Monitors, Desktop Computers, Printers and kitchen equipment. In our model, this office equipment shows up as residual term R as LightMon could not observe the state of those devices. The power supply

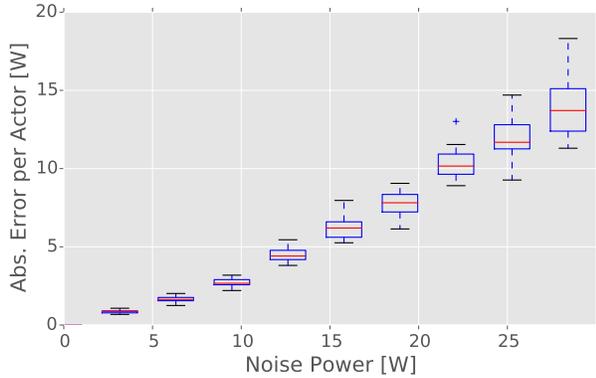


Figure 9: Error dependent on noise level. The simulated noise is based on logarithmic series distribution ($a = 0.85$) and scaled with the denoted power.

in our test environment is realized by two independent sub-distributions with three phases each. All six phases were monitored using two EMU Professional series electricity meters. Our LightMon system obtains power readings via MODBUS interface with a sample rate of 0.5Hz. The building is equipped with a KNX-bus system for controlling lights. The LightMon connects to the KNX-bus using a commercially available ABB KNX to IP gateway. Although all four floors of the building share the same KNX bus, our system also received events from other parts of the building.

The building department of our university provided us with a mapping of KNX actuator addresses to the rooms where the actuators are installed. We used this mapping to replace the KNX actuator address with its symbolic name for better comprehensibility. If no such mapping is available for the building, the live monitoring feature of LightMon could be used to generate such a symbol table. In this case, one has to manually switch all lights and record the corresponding BCN addresses using live monitoring. We use this approach to check the correctness of the provided symbol table.

In our test environment five different kinds of lights are switched by the BCN. More precisely, in our environment 29 Desktop lights (S), 8 Wall lights (W), 26 Ceiling lights (D), and 5 Bathroom lights are installed. We searched for lights not attached to the BCN and found one light which was privately purchased by an employee. The rest of the electricity consumption is caused by ordinary office equipment, printers, and kitchen equipment.

Ground Truth Estimation: In order to obtain ground truth data for all lamps present in the environment, we conducted a measurement when the whole floor was unoccupied. During that measurement, all disengageable loads were turned off to avoid random variations in the power draw caused by non-lighting equipment. Then, we switched each lamp five times on and off. During that time, we monitored the power consumption on all phases. The ground-truth data also shows the deviation from our assumed power model until the lamps reach their steady state power consumption. However, we manually reviewed the collected ground data for each lamp to ensure the absence of noise effects from other electric devices.

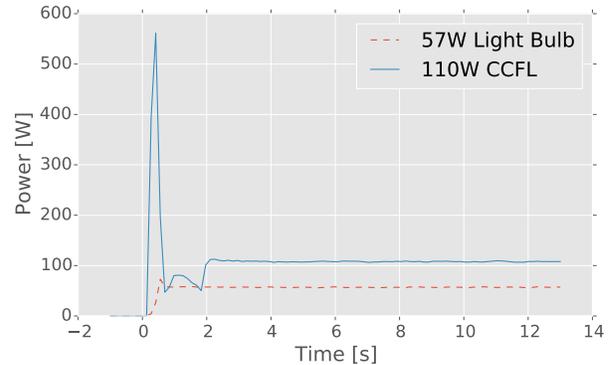


Figure 10: Turn-on transient for commonly used CCFL and halogen lamp.

Manual Electricity Audit: In addition to that, we manually counted the number of installed light bulbs and read manufacturer information printed on each installed light bulb. The manufacturer information differs from our ground truth measurement by approximately 4.92%. A comparison on individual light level showed that the ground truth measurement is almost 1W higher than it is stated by the manufacturer information. The biggest single consumer was the two independent lights installed in our hallway. Each of them consists of 28 light bulbs with 16W power consumption per light bulb. Both lights are controlled by a movement sensor which switches those lights for 15 minutes if movement is detected.

In the whole test environment, we found six lights which were not measurable by LightMon as they are connected to a separate sub-distribution. Those lights are installed in the rest room areas. Five lights were broken during the ground truth measurement and for 17 BCN switching actuators, no corresponding light was connected. All those 17 places were not equipped with a desk. Two lights even had a programming error which caused an inverted reaction to switching command, e.g. they switch on when receiving a BCN switch off event. We received valid, non-zero measurement for 39 individual lights.

We compared the power consumption matrix derived from our ground truth measurement with a power consumption matrix obtained from LightMon after observing the office environment for 30 days during a sunny summer month. The average absolute error between ground truth measurement and observation is 11.34%. However, this error is not evenly distributed. All 13 lights which were frequently used have a very low absolute error. The rest of the lights were rarely used and show significant estimation errors. While acceptable for most applications, it may be that the reported observation error is too high for particular applications. However, in such cases, LightMon could be configured to automatically conduct a ground truth measurement at night times or on weekends.

5 Application Scenarios

In this section we describe the findings of one year LightMon usage in one floor of our institutes building. The Light-

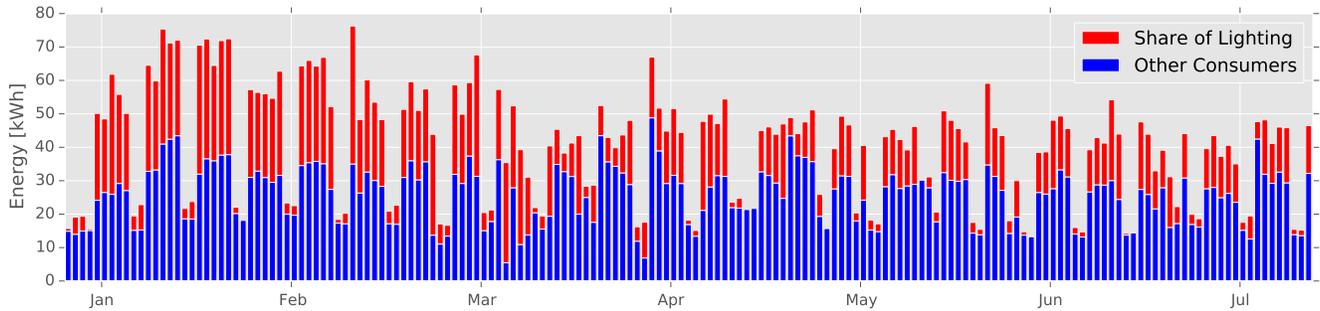


Figure 11: Share of Lighting on the buildings total electricity consumption.

Mon system collected $n=356$ days of electricity consumption together with corresponding BAN light switching events from July 2014 till July 2015. The whole data set is available online¹. The data was recorded continuously except for a period of 27 days in November and December when our system was not operating properly.

We obtained a power consumption matrix from this data and calculated the electricity consumption for each light individually. As result we get a multidimensional time series stating the electricity consumption of each light for each point in time. In order to make this large matrix more comprehensible for people, we aggregated the data in three different views. Those views show the electricity consumption for lighting as daily profile, the share of lighting on the total electricity demand varying between different month as well as the electricity used for lighting different rooms and spaces of the building. Thus, the electricity consumption has a high correlation with the total electricity consumption. Furthermore, the lighting profile reflects the usual working conditions at our department. The first employee starts at 7:15 in the morning and the last employees leave at round about nine in the evening. We calculated the share of lighting on the total electricity consumption for each day from Jan till July to get a better understanding of its influence on the total electricity demand. The result of this analysis is shown in Figure 11. As expected, the results show a strong seasonal behavior. During the winter months (Jan, Feb), the share of lighting is in total as high as 42%. In spring and summer the share lighting on the total electricity consumption drops down to a total percentage of 28%. In the period from January to July our monitored environment consumed 6.8MWh of electricity from which 2.4MWh (35%) was spent for lighting. When interpreting Figure 11 it is important to note that this graph is not representative for the big stock of available buildings, as other buildings are equipped with other electricity consumers, have different heating or cooling concepts and the inhabitants have different behavior. For future work, it would be interesting to compare the share as well as the total amount of lighting on the electricity consumption for different kinds of buildings. To answer the question why such a high share of electricity is used for lighting, we analyzed the electricity consumption on room level. For that purpose, we queried LightMon to a

list of the installation location of all lights, its type (c.f. Section 4.6), its average daily electricity consumption, its daily cumulated switch on time, and the average cycle length between a switch on command and a switch off command. This automatically generated summary table is useful for selecting the most efficient lighting solution for differently used rooms. Depending on the total switch on duration as well as the switching frequency, different lamps could be selected. We found that in our observed building, the hallways dominate the electricity demand for lighting. Therefore, investments in energy efficiency should focus on the hallways first. In each hallway 28 light bulbs with 16W each are installed. The lighting of the hallway is controlled by movement sensors which switch on the lighting for 15 minutes if movement is detected. This causes a constant illumination of our hallways for over 8 hours every day. However, the movement sensors do not consider if the hallways are already sufficiently lighted by the sun. Therefore, the single biggest saving potential could be realized by adding indoor luminance sensors in our hallways and using them to avoid over-provisioning. Alternatively, the currently installed light bulbs could be replaced with more energy efficient solid state lights.

6 Discussion

The design of a specific system for monitoring the electricity cost of lighting results in unique technical as well as non-technical design decisions. Those are discussed with much broader context in the next section.

Switching in un-occupied situations: Instead of live-monitoring the changes in the power signal of a building, the system could rely on one-time calibration. With other words, the system could automatically switch all loads during night time to determine the power draw of each actuator individually. While being much simpler to implement and evaluate, this active approach has severe disadvantages. First, switching all loads causes additional and unnecessary electricity demand. Second, online detection of faults becomes impossible. And third, the auto-calibration at night times seems strange for human spectators being currently in the environment. Therefore, we use this approach only as reference measurement for verification purposes.

Manual Energy Audit: Instead of installing additional hardware for monitoring the electricity consumption of lighting, one could conduct a manual audit of all present lights. This step would involve enumerating all lights present in a

¹Download available on the LightMon Website at <https://nglrt.github.io/LightMon/>

building together with their power rating and estimating the average on-time for each light individually. This work is labor intensive and most probably much more expensive than installing LightMon. However, we executed this procedure once in our test environment in order to verify ground truth.

Applicability in legacy buildings: The proposed system relies on the existence of a building wide BAN for switching the lights. Thus, the system is not easily applicable in legacy buildings. The installation of additional luminance sensors or the usage of Smartphone-integrated sensors could mitigate this fact. In the long term, this work around is not necessary as the prices for networked lighting drops rapidly. With affordable networked lights, each lamp in a building will be able to provide the necessary state information.

7 Conclusion

We presented the LightMon system, which is capable of measuring the electricity consumption of individual lights by correlating switching events with their influence on the total power consumption of a building. Our proposed system is seamlessly deployable in modern buildings equipped with a building automation network. It estimates power consumption on the level of individual lights with an error as low as 11.34% and provides a web based user interface for live consumption monitoring as well as viewing historic electricity consumption of the lighting installation. LightMon adapts itself to the characteristics of any building where it is installed and requires neither additional explicit training nor manual inputs. By using our LightMon system, it becomes possible to track the financial cost of electric lighting in a building. This is the key information required to implement energy savings projects on an economical basis.

Acknowledgements

This work is co-funded by the Social Link Project within the Loewe Program of Excellence in Research, Hessen, Germany.

8 References

- [1] A. Alhamoud, F. Ruettiger, A. Reinhardt, D. Burgstahler, D. Böhnstedt, and R. Steinmetz. SMARTENERGY . KOM : An Intelligent System for Energy Saving in Smart Home. In *LCN'14*, 2014.
- [2] A. Alhamoud, P. Xu, A. Reinhardt, F. Englert, P. Scholl, D. Boehnstedt, and R. Steinmetz. Extracting Human Behavior Patterns from Appliance-level Power Consumption Data. In *EWSN'15*, 2015.
- [3] S. Barker, S. Kalra, and D. Irwin. PowerPlay : Creating Virtual Power Meters through Online Load Tracking. In *BuildSys'14*, 2014.
- [4] N. Batra, J. Kelly, O. Parson, H. Dutta, W. Knottenbelt, A. Rogers, A. Singh, and M. Srivastava. NILMTK : An Open Source Toolkit for Non-intrusive Load Monitoring. In *ACM e-Energy'14*, 2014.
- [5] C. Beckel, W. Kleiminger, R. Cicchetti, T. Staake, and S. Santini. The eco data set and the performance of non-intrusive load monitoring algorithms. In *ACM BuildSys'14*, Memphis, 2014.
- [6] C. Beckel, W. Kleiminger, T. Staake, and S. Santini. Improving Device-level Electricity Consumption Breakdowns in Private Households Using ON/OFF Events. *CONET'12*, 2012.
- [7] P. Bertoldi and B. Atanasiu. Electricity Consumption and Efficiency Trends in European Union. 2009.
- [8] M. Clark, B. Campbell, and P. Dutta. Deltaflow: Submetering by Synthesizing Uncalibrated Pulse Sensor Streams. In *ACM e-Energy'14*, 2014.
- [9] J. Conti. Annual Energy Outlook. *Energy Information Administration (EIA)*, 2015.
- [10] S. Darby. The EffectiveNess of Feedback on Energy Consumption. *A Review for DEFRA of the Literature on Metering, Billing and direct Displays*, 2006.
- [11] D. Egarter and W. Elmenreich. Autonomous Load Disaggregation Approach based on Active Power Measurements. In *PerEnergy'15*, 2015.
- [12] F. Englert, I. Diaconita, A. Reinhardt, A. Alhamoud, and R. Steinmetz. Reduce the number of sensors: Sensing acoustic emissions to estimate appliance energy usage. In *BuildSys'13*, 2013.
- [13] F. Englert, P. Lieser, A. Alhamoud, and R. Steinmetz. Electricity-Metering in a Connected World: Virtual Sensors for Estimating the Electricity Consumption of IoT Appliances. In *FiCloud'15*, 2015.
- [14] F. Englert, T. Schmitt, S. Kössler, A. Reinhardt, and R. Steinmetz. How to Auto-Configure your Smart home? High-Resolution Power Measurements to the Rescue. In *e-Energy'13*, 2013.
- [15] C. Fischer. Feedback on household electricity consumption: a tool for saving energy? *Energy Efficiency*, (1), May 2008.
- [16] S. Gupta, M. S. Reynolds, and S. N. Patel. ElectriSense: single-point sensing using EMI for electrical event detection and classification in the home. *Ubicomp '10*. ACM, 2010.
- [17] S. Hay. *A model personal energy meter*. PhD thesis, University of Cambridge, 2011.
- [18] D. Irwin, S. Barker, A. Mishra, and P. Shenoy. Exploiting home automation protocols for load monitoring in smart buildings. In *BuildSys'11*, 2011.
- [19] X. Jiang, S. Dawson-Haggerty, P. Dutta, and D. Culler. Design and Implementation of a High-Fidelity AC Metering Network. In *ACM IPSN'09*, 2009.
- [20] D. Jung and A. Savvides. Estimating building consumption breakdowns using on/off state sensing and incremental sub-meter deployment. In *SenSys'10*. ACM, 2010.
- [21] J. Kelly and W. Knottenbelt. The UK-DALE dataset, domestic appliance-level electricity demand and whole-house demand from five UK homes. *Scientific Data*, 2, 2015.
- [22] Y. Kim and T. Schmid. ViridiScope: design and implementation of a fine grained power monitoring system for homes. In *ACM UbiComp'09*, 2009.
- [23] J. Kolter and M. Johnson. REDD: A public data set for energy disaggregation research. In *ACM SIGKDD'11*, San Diego, 2011.
- [24] N. Pathak, M. Khan, and N. Roy. Acoustic based appliance state identifications for fine-grained energy analytics. In *PerCom'15*, 2015.
- [25] A. Peruffo, A. Pandharipande, D. Caicedo, and L. Schenato. Lighting control with distributed wireless sensing and actuation for daylight and occupancy adaptation. *Energy and Buildings*, 97, June 2015.
- [26] B. Roisin, M. Bodart, A. Deneyer, and P. D'Herdt. Lighting energy savings in offices using different control systems and their real consumption. *Energy and Buildings*, 40(4), Jan. 2008.
- [27] M. Saha, S. Thakur, A. Singh, and Y. Agarwal. EnergyLens : Combining Smartphones with Electricity Meter for Accurate Activity Detection and User Annotation. In *ACM e-Energy'14*, 2014.
- [28] E. Tetri and P. Bhusal. Annex 45 guidebook on energy efficient electric lighting for buildings. 2010.
- [29] A. Uttama Nambi, A. Reyes Lua, and V. Prasad. LocED: Location-aware Energy Disaggregation Framework. In *BuildSys '15*, pages 45–54. ACM, 2015.