

HeatProbe: A Thermal-based Power Meter for Accounting Disaggregated Electricity Usage

Bo-Jhang Ho¹, Hsin-Liu (Cindy) Kao¹, Nan-Chen Chen¹, Chuang-Wen You²,
Hao-Hua Chu¹, Ming-Syan Chen²

¹Department of Computer Science and Information Engineering, National Taiwan University
1, Sec. 4, Roosevelt Rd., Taipei 106, Taiwan
{b96118, b95701241, b97006, hchu}@csie.ntu.edu.tw

²Research Center for Information Technology Innovation, Academia Sinica
128, Sec. 2, Academia Rd., Taipei 115, Taiwan
{cwyu, mschen}@citi.sinica.edu.tw

ABSTRACT

To promote energy-saving behavior, disaggregating electricity usage is critical for increasing consumer awareness of energy usage behavior. This study proposes HeatProbe, a thermal-based power meter system that uses thermal imaging to track disaggregated appliance usage. We have designed, prototyped, and tested the HeatProbe system. Results show that HeatProbe successfully senses individual appliance operating durations with an average error of 125.03 seconds, achieving 80.2% appliance power accounting accuracy in different appliance usage scenarios.

Author Keywords

Power Consumption Monitoring, Per-appliance Power Disaggregation.

ACM Classification Keywords

H5.m. Information interfaces and presentation (e.g., HCI): Miscellaneous.

General Terms

Design, Experimentation.

INTRODUCTION

Energy conservation has recently become a significant concern [1], and has long been an important goal of UbiComp research on promoting energy consumption awareness [2, 3, 4]. Since utility bills aggregate electricity usage, it is difficult to determine disaggregated energy consumption. If individuals can be made aware of their major energy consuming behaviors (e.g., the type and operating duration of the consuming appliance), they can explore ways to make positive changes in their energy usage. As emphasized by Stern [17], this information awareness on appliance usage behavior and incentives for changing environmentally significant consumer behavior can have synergistic effects to conserve energy.

This study proposes the HeatProbe system to track the energy consumed by individual appliances. The proposed system disaggregates total power usage from a master power meter into individual appliance power usage by tracking the heat patterns generated by running appliances. This approach is based on indirect sensing. Rather than directly sensing the power line connecting to each appliance, HeatProbe uses a thermal camera and thermal imaging algorithms to recognize heat patterns from appliance usage. Figure 1 shows heatmaps of appliances captured at two different time points. The upper heatmap (Fig. 1(a)) shows the appliance surface near room temperature in the power-off state. The lower heatmap (Fig. 1(b)) shows that a user has entered the space and turned on various appliances whose surface temperature exceeds room temperature. Since appliances produce heat as a byproduct of their operation, HeatProbe determines changes in an appliance's power state by tracking its surface temperature. A rising surface temperature implies that the appliance has been turned on, while a falling surface temperature implies the appliance has been turned off.

Several recent power metering systems, such as ElectricSense [2], Viridiscopes [3], and contactless EMF sensing method [4], allow the tracking of *per-appliance* electricity consumption by analyzing the electro-magnetic signals of each appliance in the home. However, these systems require per-home calibration or a training process to determine the parameters (e.g., calibrated parameters for estimating power consumption or electro-magnetic fingerprints) of different appliances. Recognizing the difficulty of these calibration/training efforts, this study proposes the HeatProbe system. This system aims to automate the tracking of disaggregated appliance energy usage footprints by indirectly sensing the runtime temperature-changing trends on appliance surfaces. Since all electric appliances generate heat during operation, individual appliance usage can be inferred from the detected appliance on/off events.

The main contribution of this study is its novel thermal approach to disaggregate whole electric energy usage into appliance-level energy consumption. This study presents thermal imaging algorithms to recognize heat patterns

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

UbiComp '11, September 17–21, 2011, Beijing, China.

Copyright 2011 ACM 978-1-4503-0630-0/11/09...\$10.00.


from appliance surfaces and human body surfaces, thus accounting for individual appliance electricity usage. Experimental results showed that the proposed thermal approach sensed individual appliance operating durations with an average error of 125.03 seconds and achieved 80% appliance power accounting accuracy in different appliance usage scenarios.

RELATED WORK


Existing studies on monitoring electric energy consumption

Commercial electricity monitoring solutions, such as Cent-a-Meter [5] and TED [6], provide aggregated whole-home electricity usage with a single-point inline power meter. These solutions produce accurate aggregated power monitoring results. Jiang *et al.* [7] designed a wireless sensor network to monitor AC power usage with customized *ACme* nodes. All nodes in this system route the power readings back to a server over the IP network. All collected power data is then processed on the server and can be visualized by users, who view their individual real-time energy consumption using a Web interface. Considering the installation difficulty for end-users, Patel *et al.* [16] designed an easily deployable power meter that consists of a sensor unit with wireless radio interface attached to the electrical breaker panel and a calibration device plugged in to each electrical outlet. This system computes real-time power consumption by sensing the anisotropic magnetoresistive (AMR) effect caused by changing magnetic fields and collecting information from the plug-in devices on outlets.

Many researchers have investigated various ways to monitor appliance-level power consumption. By analyzing power readings from in-line power meters, Non-Intrusive Load Monitoring (NILM) [8, 9] identifies the sudden changes in the voltage or current readings caused by the state changes of appliances. Based on these detected appliance on/off events, NILM disaggregates whole-home power consumption into appliance-level components. Rowe *et al.* [4] exploited an electromagnetic field (EMF) sensor to assist the training phase of NILM. However, a machine learning procedure is required to learn models for specific appliances. ViridiScope [3] proposed an indirect-sensing approach to estimate current power consumption based on the magnitude of the sensed magnetic field near an appliance. Intermediate appliance state changes can also be detected using light and acoustic sensors. ElectriSense [2] enables the tracking of different appliance usages by analyzing the continuously high frequency electromagnetic interference (EMI) signals generated when an appliance is running. By extracting the features embedded in the emitted EMI signals, ElectriSense can classify the usage of different appliances using a machine learning toolkit. In contrast, HeatProbe captures the runtime heatmaps using a thermal camera to segment appliance heat and detect significant heat increases/decreases on the appliance surfaces. Using this data, the system infers the appliance operating duration



(a) Before turning on appliances



(b) After turning on all appliances

Figure 1. Two frames captured (a) before and (b) after all appliances are turned on. These five appliances are a lamp, a monitor, and oven, a PC and a printer.

and power usage without performing calibration/training steps.

Prior Work Using Thermal Sensing

Previous researchers have explored various possibilities on applying thermal imaging in the HCI research. StressCam [10] presents a method to detect people’s stress level using thermal cameras. Based on the correlation between people’s stress level and their forehead temperature, StressCam measure their physiological stress based on the temperature of their foreheads. Yun *et al.* [11] extended this thermal imaging method to observe players’ physiological stress level while playing games. Based on the detected stress level, the system can infer the difficulty level in a game faced by individual players and then adjust the difficulty of the game dynamically based on players’ skill levels. Heat-Wave [12] explores thermal sensing to enhance user surface interaction. Using a thermal camera to track residential heat traces left on an interactive surface, their system detects various finger gestures such as hovering, touching, target selection, etc.

THEORY OF OPERATION

The HeatProbe system infers per-user power consumption using one in-line power meter and one thermal camera. According to the electric circuit theory [15], the rate of electrical energy dissipation (P) due to resistance is proportional to the current (I) and resistance (R) and based on the following equation:

$$P = I^2R \quad (1)$$




Figure 2. Resulting average surface temperature of two exemplar appliances (a toaster and a monitor) when the appliance is turned on/off. These two cases show that surface temperature increases (decreases) when the appliance turned on (off).

Dissipated energy due to the internal resistance of an appliance is converted into the form of thermal energy, i.e., heat, which causes the temperature on the appliance surface to rise.

HeatProbe uses a thermal camera to capture this temperature change on appliance surface. Figure 1 shows the temperature difference between two heatmaps captured at two time points. The upper heatmap (Fig. 1(a)) shows the appliance surface near room temperature in the power-off state. The lower heatmap (Fig. 1(b)) shows that a user has entered the space and turned on appliances, including a desk lamp, an LCD monitor, a printer, a PC, and an oven. The surface temperature of these appliances rises, exceeding room temperature. HeatProbe uses the change in the appliance surface temperature to infer the appliance on/off state. This heat-producing phenomenon is universal to almost all electric-powered appliances.

Assumptions and Limitations

The current HeatProbe system works based on the following assumptions. (1) The system tracks electricity usage from appliances that are plugged into the building circuit such that the inline master power meter can sense their electricity usage. For example, the system cannot track electricity consumed by mobile devices that run on their own batteries. (2) The system cannot track an appliance whose surface is not within view of the thermal camera, e.g., appliances that are hidden underneath a desk. (3) The system recognizes only appliance on/off binary state. However, some appliances, such as refrigerators, have multiple power

states. Since an intermediate power state change often produces a small temperature change, the current thermal camera cannot detect heat-changing events caused by the intermediate appliance state change. Thus, the current system implementation assumes that the appliance consumes constant amount of power. However, most shared appliances have limited operating states and tend to be used for a limited amount of time. Thus, the amount of under- (or over-) estimated power consumption should be relatively insignificant over a longer monitoring period.

SYSTEM DESIGN

Figure 3 presents the HeatProbe system design. HeatProbe implements five inference modules to derive appliance power consumption from the raw heatmaps and power readings: (1) heatmap segmentation, (2) power event detection, (3) appliance usage detection, (4) appliance operating duration resolution, and (5) appliance power accounting. This section first provides a brief description of how these five inference modules operate in the HeatProbe system.

(Step 1). *Heatmap segmentation* processes each heatmap and partitions it into multiple segments, each of which corresponds to either a heated (running) appliance or a human body. For example, after segmenting the heatmap in Fig. 1(b), HeatProbe locates six heat segments corresponding to five appliances and one human body. HeatProbe then computes the surface temperature of each appliance segment. By continuously computing the surface temperature of each appliance segment over time, the system determines temperature changes on the appliance surface. Rising (falling) temperature indicates a *heat event* with an on (off) appliance power state change. HeatProbe also identifies and tracks the human bodies by continuously analyzing heatmaps.

(Step 2). *Power event detection* analyzes aggregated power readings obtained from the master power meter and identifies *power events* correlated to changes in an appliance's on/off power state. For example, Fig. 2(a) shows




Figure 3. HeatProbe system design.




Figure 4. Steps in the heatmap image segmentation module.

the aggregated power readings including two pairs of power-on and -off events. They correspond to (e1) power-on of a toaster at the 16th minute, (e2) power-off of the toaster at the 17th minute, (e3) power-on of an LCD monitor at the 20th minute, and (e4) power-off of the LCD monitor at the 56th minute.

(Step 3). *Appliance usage detection* monitors any significant temperature increases/decreases for each appliance segment. Because electric energy dissipates on the internal resistance as a form of thermal heat, the generated heat will cause a noticeable temperature increase on the surface of an appliance. For example, Fig. 2 shows significant temperature increases occurring at the 16th and 21th minutes. After the toaster (monitor) turns off, the surface temperature drops back to the room temperature, as indicated at the 22th (56th) minute in Fig. 2 (b) and (c). This allows the system to infer an appliance on/off state change by detecting significant temperature increases/decreases in an appliance segment.

(Step 4). *Appliance operating duration resolution* captures the relationship between an appliance’s power-on event (the cause) and its surface temperature increase event (the effect). Similarly, this step captures the relationship between an appliance’s power-off event (the cause) and the subsequently temperature decrease event (the effect). For example, Fig. 2 shows the matching of four pairs of power-on (-off) and heat-rising (-falling) events. The first pair occurs around the 16th minute, when the user switches on the toaster. This action produces a power-on event at the 16th minute, and the heat-rising event from the image segment of the toaster from the 16th to the 19th minute. Three other event pairs occur at the 21th, 22th, and 56th minutes. The system determines the starting/ending times of each appliance’s usage by matching and linking these events.

(Step 5). *Appliance power accounting* estimates the average power consumption of each appliance based on the power changing amounts of power-on (off) events. For example, Fig. 2 (a) shows pairs of power-on and power-off events for a toaster. These events correspond to (e1) power-on of a toaster at the 16th minute (a 600 Watt power in-

crease), (e2) power-off of the toaster at the 17th minute (a 600 Watt power decrease). The appliance power consumption can be approximated by the amount of energy consumed in an average power (600 Watt) over the entire appliance usage session.

The following sections describe each module in details as follows.

In-line Power Meter and Thermal Camera

The system uses two sensors: an in-line power meter and a long-wavelength infrared (LW-IR) thermal camera. Existing power meters installed by the utility providers can provide data on aggregated power consumption.

The infrared/thermal camera measures the quantity of infrared (far-infrared, F-IR, or long-wavelength infrared, LW-IR) radiation emitted by any object. Based on Planck’s law, the electromagnetic radiation from a room-temperature object falls mainly in the IR spectrum, with the radiation intensity directly proportional to the fourth power of the object’s absolute temperature. By measuring the radiation intensity emitted in the IR spectrum, the thermal camera can passively generate a heatmap of an object in the view range without any illumination sources. This approach is unlike the *reflected* infrared operating in the near-infrared (N-IR) spectrum.

Heatmap Segmentation

Figure 4 shows a sequential block diagram of the heatmap segmentation module, which includes the following steps: (1) human body location tracking, and (2) appliance heat pixel labeling steps. These two steps are described as follows.

Human body location tracking: Given raw heatmaps, the human body location tracking component finds moving heating sources (i.e., humans) by identifying pixels with sudden temperature changes. Because the temperature values of human bodies are significantly higher than room temperature, the movement of a human body should cause abrupt temperature increases/decreases on pixels in or near human bodies between two time-adjacent frames. By identifying pixels with abrupt temperature changes, it is possible to segment out a body area by drawing a bounding contour for a clustered region of pixels with sudden temperature changes. The same procedure can be repeated for all time-adjacent heatmaps to find a corresponding human mask, which labels all pixels in bounding contours as human-type pixels. Subsequent modules can use this human mask to retrieve the location of a user, which is given by the coordinates of all pixels within the selected human body area.

Appliance heat pixel labeling: To locate the regions of running appliances, the system uses the human-suppressed heatmaps to find heated appliance pixels with temperatures above the room temperature (24.3 °C). First, the system dynamically selects a cut-off temperature in each heatmap to segment out appliance heat pixels. The system then analyzes the temperature distribution of pixels in each heatmap. Figure 5 shows an example temperature distribution of a




Figure 5. Temperature distribution of all pixels in a frame. The peak (red dotted line) is fit by the Gaussian distribution, with a mean value of μ and a standard deviation of σ .

heatmap. Due to the thermal camera's limited temperature sensing range ($23\text{ }^{\circ}\text{C} \sim 35\text{ }^{\circ}\text{C}$), all temperature values are limited within the sensing range. All pixels with temperatures higher (lower) than maximum (minimum) sensing temperature are set to the maximum (minimum) value. This results in large number of pixels at the minimum ($23\text{ }^{\circ}\text{C}$) and maximum ($35\text{ }^{\circ}\text{C}$) temperature. This temperature distribution also has a peak located around the room temperature ($24.3\text{ }^{\circ}\text{C}$). The whole distribution is fit by a Gaussian function to extract the mean (μ) and variance (σ^2) parameters shown as the red dotted line in Fig. 5. The cut-off temperature is set as the mean value (around the room temperature) plus one standard deviation, $\mu + \sigma$. After determining the cut-off temperature for a heatmap, it is possible to label all pixels with temperature values greater than the cut-off temperature as appliance-type pixels in a corresponding labeling frame, as indicated in the middle row of Fig. 6. Repeating the same procedure for each raw heatmap identifies the corresponding appliance heat pixels. The resulting appliance heat pixels are then used to generate appliance masks for detecting the appliance usages.

Power Usage Detection

The raw power readings in Fig. 2(a) are measured by an in-line power meter that samples the power consumption of appliances once per second. These raw power readings can be smoothed using a median filter with a sliding window of five samples. The starting (or ending) time of a power-on (or -off) event can be determined by finding a significant change in the value of the smoothed power curve, which is caused by turning on (off) an electric load. A significant change of a power event of the i^{th} appliance usage session, defined as Δp_i , gives the power demand of that appliance. For example, Fig. 2 (a) shows that the power readings increase (decrease) at the 20th (56th) minute when the LCD monitor is turned on (off), with an increasing (decreasing) amount of 25 (25) Watts. Therefore, a power event is defined as power changing vectors $P_i = (t_i, \Delta p_i)$, where t_i , and Δp_i are the starting (ending) time, and the increasing (decreasing) amount of the power readings of the power on (off) event of the i^{th} appliance usage session, respectively.

Appliance Usage Detection

This module involves the following four steps: (1) temperature-based spatial smoothing, (2) appliance heat mask gen-




Figure 6. Appliance mask generation steps. The top row includes raw heatmaps captured in five adjacent times. The middle and lowest rows represent the corresponding labeling frames and masks. The right-most frame contains the resulting appliance heat areas for a lamp and a monitor.

eration, (3) heat-rising event detection, and (4) heat-falling event detection.

Temperature-based spatial smoothing: The temperature-based spatial smoothing step first performs noise reduction on the observed temperature. Each frame is smoothed by finding the average temperature of each 3×3 neighborhood of pixels. The same procedure is repeated for all pixels in a heatmap, generating a smoothed heatmap. These smoothed heatmaps are processed in subsequent steps to detect heat events.

Appliance heat mask generation: This step obtains the appliance heat masks corresponding to different appliance usage sessions, which indicate the pixel locations of running appliances. During appliance operating, the number of heat pixels (i.e., whose temperature is above the room temperature) of an appliance might vary due to temperature increases (decreases) on the appliance surface or appliance form changes caused by usage behaviors (e.g., opening/closing the LCD monitor of a laptop). To locate an appliance's pixels, the system joins all heat pixels during an appliance operating session to form a union area, called an *appliance heat area*. This appliance heat area can effectively monitor the temperature-changing trends during the appliance operating. This area appears in the rightmost frame of the middle row of Fig. 6. An appliance mask for a specific time can be generated by unmasking pixels in appliance heat areas for all appliances running at that time. Repeating the same procedure for each time instance produces a series of time-ordered appliance masks, as shown in the lowest row of Fig. 6. The subsequent steps of this module use this appliance mask series to retrieve the coordinates of all pixels within each appliance.

Heat-rising event detection: This step analyzes the smoothed heatmaps to detect the heat-rising events of all appliances. Turning on an appliance should persistently dissipate electric energy, and thus cause a noticeable temperature increase on the appliance surface. The heat-rising process shows a slow but persistent temperature increase. Thus, the proposed system takes the average temperatures




Figure 7. The average temperature for an appliance heat area. The blue line is the original data curve, and the red line is the smoothed curve after applying a low-passed filter.

of all pixels for each appliance (i.e., all pixels in each appliance heat area) every second, producing the temperature trend depicted in Fig. 7. To detect the rising trends, the raw average temperature is smoothed using a low-pass filter (the red line in Fig. 7). The smoothed signal is then analyzed with a sliding window (Δt_{on}) of 60 samples or 1 min of sensed average temperature.

The occurrence of a heat-rising event corresponds to a condition in which the temperature difference between the beginning and ending time in the sliding window exceeds a threshold $\Delta T = 0.2^\circ\text{C}$, indicating a slow but substantial change. This is determined by the average heat-rising rate observed from the empirical results of the appliances in our experiment. After the starting time of an event, the temperature changing rate (r_i) of an event of the i^{th} appliance usage session is equal to the maximum slope in the sliding window. This value represents the heat generation characteristic of the appliance when turned on.

Heat-falling event detection: The proposed system detects the occurrence of a heat-falling event by monitoring the major temperature-changing trend of all pixels. Due to the thermal conductivity of the materials surrounding the appliances (e.g., the air, the surface of the desk, etc.), it takes some time to disperse the generated heat left inside an appliance. This makes the surface temperature fall at a slow rate after an appliance is turned off. The heat-falling rate is not uniform across the appliance surface. To detect the temperature changing trend of each pixel, the system calculates the temperature difference (T_{diff}) of a pixel between the current temperature and the previous temperature recorded 60 seconds earlier (a time interval long enough to reflect the falling trend). The system can then classify each pixel into one of three temperature-changing categories: a temperature-rising pixel ($T_{diff} > T$), a temperature-falling pixel ($T_{diff} < -T$), and a temperature-fixed pixel ($-T \leq T_{diff} \leq T$). The temperature threshold, T , is determined by the temperature resolution of the thermal camera ($T=0.3^\circ\text{C}$ in the current implementation). A heat-falling event occurs when there are more temperature-falling pixels than temperature-rising pixels in an appliance heat area. A trend indicator (TD) calculates the difference between the percentage of temperature-rising pixels and the percentage of temperature-falling




Figure 8. The trend indicator for an appliance heat area. The blue line is the original data curve, and the red line is the smoothed curve after applying a low-passed filter.

pixels in an appliance area. Fig. 8 shows when an appliance is turned off, TD drops below zero because the number of temperature-falling pixels exceeds the number of temperature-rising pixels. After the surface temperature of the appliance returns to the room temperature, TD also goes near zero. The system calculates TD for each appliance heat area in every second interval (Fig. 8). A low-pass filter smooths the TD samples (the red line in Fig. 8) with a sliding window (Δt_{off}) of 90 TD samples over one and a half minutes. An appliance-off event corresponds to a drop in TD below zero in the sliding window exceeding a threshold (ΔTD) of 10%. This means the percentage of temperature-falling pixels exceeds the percentage of temperature-rising pixels by 10%.

Appliance Operating Duration Resolution

This module resolves the appliance-on (-off) time and operating duration of each appliance usage session. Each appliance-on (-off) state change generates a pair of power and heat events that are correlated with each other in event occurrence time. Therefore, this module infers the appliance-on (-off) times by finding the optimal matched pairs of power and heat events based on time proximity, which can be formulated as an event assignment problem [13]. Each operating duration can then be obtained by calculating the time difference between on and off time. This module involves the following three steps: (1) appliance-on time, (2) appliance-off time, and (3) operating duration resolution steps.

Appliance-on time resolution: The appliance-on time is determined by power-on and heat-rising events matching algorithm. Appliances requiring a pre-heating step in preparation for further action (e.g., printing papers) tend to consume a larger amount of power (i.e., a larger Δp value). This pre-heating step also causes the appliance surface temperature to increase rapidly (i.e., at a higher heat-rising rate r_i) when entering the full operating mode, and thus instantly triggers a significant temperature increase. Based on this observation, the initial matching stage selects power events with a large Δp value and thermal events with a high heat-rising rate. To identify appliances with high power consumption, the system computes the average power-changing amount from all power-on events and then identify high-power events whose Δp is above the average value.




Figure 9. An example power-heat matching graph. In (a), three power events (P₁, P₂, and P₃) are connected with three heat events (H₁, H₂, and H₃) with edges with costs listed in (b).

Similarly, the system designates heat-rising events with temperature changing rates above the average rate. After obtaining these high-power (-heat) events, the system performs matching to join heat and power events. Figure 9 shows an example in which high-power and high-heat events are represented as two types of vertices: power and heat vertices. A heat vertex represents an i^{th} heat event ($H_i^{\text{on}}(t_i^{\text{on}}, r_i^{\text{on}})$) occurring on t_i^{on} with a heat-rising rate r_i^{on} . A power-type vertex representing a j^{th} power event ($P_j^{\text{on}}(t_j^{\text{on}}, \Delta p_j^{\text{on}})$) occurring on t_j^{on} with the power-changing amount Δp_j^{on} . An edge connecting a heat vertex to a power vertex has a matching likelihood weight (or cost) set as the time difference ($|t_i^{\text{on}} - t_j^{\text{on}}|$) between the i^{th} heat and j^{th} power events. By adding these edges between heat and power vertices creates a weighted bipartite graph.

In the next step, the system applies an event assignment algorithm [13] to find the optimal event assignment in which the sum of cost from edges is minimum. After performing the high-power and -heat event matching, the same algorithm is performed to join all remaining power-on and heat-rising events. The resulting matched on-event pair of the k^{th} appliance usage session is defined as $E_k^{\text{on}}(t_k^{\text{on}}, \Delta p_k^{\text{on}})$, where t_k^{on} and Δp_k^{on} are the corresponding appliance-on time and power changing amount of this on-event pair, respectively. Because the power-on time and heat-rising time given by events of an event pair may differ, the appliance-on time t_k^{on} is set as the power-on time of the power event included by the matched event pair.

Appliance-off time resolution: After the on-event matching step, the appliance-off time is determined by matching power-off and heat-falling events using the weighted bipartite graph and event assignment algorithm described in the previous paragraph. However, the edge cost used in the off-event matching is calculated differently, as explained below. Given an on-event pair of the k^{th} appliance usage session matched by the previous step, the edge cost between a m^{th} heat-falling event, $H_m^{\text{off}}(t_m^{\text{off}}, r_m^{\text{off}})$, and a n^{th} power-off event, $P_n^{\text{off}}(t_n^{\text{off}}, \Delta p_n^{\text{off}})$, is calculated as follows:

$$cost_{mn} = w_t \times time_cost_{mn} + w_p \times power_cost_n \quad (2)$$

$$where \quad time_cost_{mn} = \frac{|t_m^{\text{off}} - t_n^{\text{off}}|}{\max_{mn} |t_m^{\text{off}} - t_n^{\text{off}}|} \times 100 \quad (3)$$

$$power_cost_n = \frac{|\Delta p_n^{\text{off}} - \Delta p_k^{\text{on}}|}{\max_n |\Delta p_n^{\text{off}} - \Delta p_k^{\text{on}}|} \times 100 \quad (4)$$

where $time_cost_{mn}$ measures the time proximity between the m^{th} heat-falling event and the n^{th} power-off even, and $power_cost_n$ reflects the correlation of the power-changing amount between the current n^{th} power-off and the value inferred from previous matched power-on events. Thus, the $time_cost_{mn}$ is proportional to the time difference between the power and heat events ($|t_m^{\text{off}} - t_n^{\text{off}}|$). The $power_cost_{mn}$ is proportional to the difference of the power-changing amount between power -on and -off events ($|\Delta p_n^{\text{off}} - \Delta p_k^{\text{on}}|$). To simultaneously consider these two sets of parameters measured in different units, both $time_cost$ and $power_cost$ are rescaled to a unified cost range of 0~100. The weights, w_t and w_p , are set as 0.8 and 0.2, respectively, to rate the time-based cost (i.e., $time_cost_{mn}$) at a higher rank than the other cost term $power_cost_n$. After the off-event matching step, the resulting matched off-event pair of the k^{th} appliance usage session is defined as $C_k^{\text{off}}(t_k^{\text{off}}, \Delta p_k^{\text{off}})$, where t_k^{off} , Δp_k^{off} are the corresponding appliance-off time and power changing amount of this off-event pair, respectively. Similarly, the inferred appliance-off time (denoted as t_k^{off}) of the matched off-event pair of the k^{th} appliance usage session is given by the time of the matched power-off event.

Operating duration resolution: The appliance operating duration (T_k) of the k^{th} appliance usage session is calculated from the time difference between on and off time instances (i.e., $t_k^{\text{off}} - t_k^{\text{on}}$). For those appliance usage session with only matched on- (off-) event pair, the system uses an off (on) time recovery scheme to further improve resolution accuracy. Most appliances are turned on by a switch on the appliance body. Although some appliances are remotely-controlled, users often interact with or touch these appliances during usage (e.g., take printed papers from the printer). Therefore, the system uses both human-appliance interaction hints and event times of un-matched thermal events to infer missing on (off) times. The system searches for the most likely human triggering action closest to a reference time indicated by the event time of a time-adjacent un-matched thermal event. The missing appliance-on (-off) time $t_k^{\text{on}} (t_k^{\text{off}})$ is given by the time instance of the triggering action that a human body area touches (or is in closest location proximity to) the appliance heat area.

Appliance Power Accounting

The appliance power accounting module estimates the average power consumption for each appliance by the matching results provided in the previous step. Based on the previous matched power-on (-off) event pair of the k^{th} appliance usage session, the average power (P_k) can be approximated by the average of power changing amount given by the on- and off-events (Δp_k^{on} and Δp_k^{off}) using the following formula.



(a) Thermal camera



(b) Power meter

Figure 10. The prototype sensor components. The thermal camera (a) is mounted on the ceiling. The power meter (b) communicates with the data receiving module wirelessly.

$$P_k = \frac{\Delta p_k^{off} + \Delta p_k^{on}}{2} \quad (5)$$

Take the monitor usage session in Fig. 2(a) as an ple: Δp_k^{on} and Δp_k^{off} are 20 Watts in the 21th minute and 20 Watts in the 56th minute, respectively. The appliance-level power consumption in the k^{th} appliance usage session can then be approximated by the electric energy consumed in an average rate of P_k over the entire session (with session duration equal to T_k).

IMPLEMENTATION DETAILS

The current HeatProbe implementation consists of (1) a NEC TH7102MX thermal camera and (2) a wireless inline power meter (Fig. 10). The camera was mounted to the ceiling to monitor a space, e.g., cubicle spaces, public rooms, or kitchen spaces. The thermal imaging signals captured by the camera were sent through an 1394 port to a backend data processing server. In the current prototype, the server records and processes sensor data. Our thermal imaging processing software is developed based on the OpenCV libraries [14].

As for power monitoring, the inline power meter transmits power readings to the server through the Zigbee radio every second. A corresponding wireless Zigbee-based data receiving module attached to the server receives power readings. The server records power readings for subsequent processing by the power usage detection module.

EVALUATION

To measure accuracy of the HeatProbe system, this study conducts experiments involving three appliance usage scenarios with multiple users.

Appliance Usage Scenarios

Table 1 lists three appliance usage scenarios. These scenarios differ on the operating appliances, the number of participants, and the number of appliance usage events. For each scenario, we scripted a sequence of appliance usage actions for the participants to perform. The following sections describe these scripted action sequences in detail. We recruited a total of six participants for these experiments. All participants were graduate students in our department. For each set of scenarios, there were several experimental rounds with different pairs of participants: there were 5, 2, and 2 rounds for scenario #1, #2, and #3, respectively. 82 different appliance usage sessions (with durations ranging from one to 60 minutes, with an average

Environment	(1) cubicle spaces	(2) room spaces	(3) kitchen spaces
Participating Appliances	2 PCs 2 monitors 2 desk lamps 1 heater 1 toaster 1 printer	2 Notebooks 2 desk lamps 1 television 1 shredder 1 bug zapper	1 water heater 1 toaster 1 electric oven 1 microwave 1 television
# of participants for each round	2 people	2 people	3 people
# of on/off events per round	20	14	18

Table 1. Experimental appliance usage scenarios for three types of environments.

length of about 22 minutes) were collected in these usage scenarios.

Appliance usage scenario #1 (office cubicles)

Two participants worked in two adjacent cubicle spaces, and a thermal camera mounted on the ceiling monitored their electricity usage. Appliances included two PCs, two LCD monitors, and two desk lamps. Additionally, one heater, one toaster, and one printer were placed between the two cubicles. The appliance usage script for participant #1 was (1) turn on/off a PC, (2) turn on/off a LCD monitor, (3) turn on/off a desk lamp, and (4) print documents from the shared printer. The appliance usage script for participant #2 was (1) turn on/off a PC, (2) turn on/off a LCD monitor, (3) a desk lamp, (4) a heater, and (5) toast bread using a toaster.

Appliance usage scenario #2 (a meeting room)

Two participants discussed their project in a meeting room where a thermal camera mounted on the ceiling monitored their electricity usage. Appliances included two notebook PCs, two desk lamps, one TV-size screen, one paper shredder, and one bug zapper. The appliance usage script for participant #1 was (1) turn on/off his/her notebook PC, (2) turning on/off a desk lamp, and (3) the TV-size screen. The appliance usage script for participant #2 was (1) turn on/off a notebook PC, (2) a desk lamp, (3) a paper shredder, and (4) a bug zapper.

Appliance usage scenario #3 (a kitchen)

Three participants operated various appliances in a kitchen where a thermal camera mounted on the ceiling monitored their electricity usage. Appliances included a water boiler, a toaster, an electric oven, a microwave, and one TV. The appliance usage script for participant #1 was (1) boil water using an electric water heater, (2) toast bread using a toaster, and (3) heat food in a microwave. The appliance usage script for participant #2 was (1) heat a microwave oven and (2) heat food in an electric oven to heat their food, and (3) watch some TV. The appliance usage script for participant #3 was (1) make toast with the toaster, and (2) heat food in an electric oven.

Evaluation Metrics

This study defines the following metrics to measure the accuracy of event detection, matching algorithms, and the appliance power accounting results.

Event Type	P-On	P-Off	P-All	H-Rising	H-Falling	H-All
Ground-truth	82	82	164	82	82	164
True positive	77	82	159	81	78	159
False positive	2	4	5	6	5	11
Average time diff (S.D.)	4.70 (7.67)	3.43 (6.79)	4.06 (7.25)	114.67 (156.59)	108.69 (125.97)	112.25 (142.37)
Accuracy	0.94	1	0.97	0.99	0.95	0.97
Precision	0.97	0.95	0.96	0.93	0.94	0.94

Table 2. Experimental results for three scenarios. P-On (or H-Rising) indicates the power-on (heat-rising) type of events. And, P-Off (or H-Falling) indicates power-off (heat-falling) type of events. P (H) -All means the overall summary of all power (heat) events. In the 5th row, the standard deviation of the time difference is indicated by values inside of the parentheses.

- *Event time detection error:* This measures the time difference between the system-detected power/heat event and the actual, ground-truth event.
- *Event detection accuracy and precision.* This measures how accurately the HeatProbe system detects power and heat events.
- *Matching accuracy.* This measures how accurately the HeatProbe system matches power events to corresponding heat events.
- *Appliance operating duration resolution error.* This measures how accurately the HeatProbe system monitors appliance operating duration. The resolution error calculates the amount of error in the HeatProbe-detected appliance operating duration and the actual ground-truth operating duration.
- *Appliance power accounting error.* This measures how accurately the HeatProbe system accounts for appliance energy consumption. The accounting error calculates the amount of error (in percentage) in the HeatProbe-detected appliance power consumption and the actual ground-truth appliance power consumption.

Results

HeatProbe-detected events were collected and compared




Figure 11. Cumulative distribution of the appliance operating duration resolution error.

Event Type	On event	Off event	All
# of Ground-truth event pairs	82	82	164
# of correctly matched event pairs	61	65	126
Accuracy	74%	79%	77%

Table 3. Summary of the event matching accuracy.

with the ground-truth events labeled by humans observing the heatmap video recordings.

Event detection accuracy and precision

Table 2 presents statistics of the HeatProbe's accuracy performance from three experimental scenarios. True (false) positives mean the number of correctly (incorrectly) detected events by the HeatProbe system. The table also shows the average time difference between the HeatProbe-detect events and the ground-truth events.

The average accuracies for detecting power-on, power-off, heat-rising, and heat-falling events are 94%, 100%, 99%, and 95%. The average precisions for detecting power-on, power-off, heat-rising, and heat-falling events are 97%, 95%, 93%, and 94%. All average detection accuracies are above 90%. The average event time detection errors for power-on, power-off, heat-rising, and heat-falling events are 4.7, 3.43, 114.67, and 108.69 seconds. The reason for the large errors in detecting times of heat events is an inherent delay between the time when an appliance is powered on/off and the time when an appliance surface exhibits a noticeable temperature change. For example, it takes a few minutes after powering on a PC to heat up its surface enough for the thermal camera to notice a temperature difference. The large standard deviations in the event time error are due to some appliances, such as a desk lamp and a heater, that increase their surface temperature more rapidly than that of other appliances.

Appliance operating duration resolution accuracy

Table 3 summarizes the number of ground-truth events, the number of correctly matched events by our matching algorithm, and the average event matching accuracy. The overall accuracy is 0.77. The matching error is caused primarily by a few seriously-delayed heat-rising (-falling) and/or false-positive events with a large time difference from the ground truth starting (ending) time. Because the current matching algorithm maps one heat-rising (-falling) event to a power-on (-off) event based on time proximity, a seriously-delayed heat-rising (-falling) event tends to match an incorrect power event (temporarily called e_i). This mismatched result might force another heat event (which should be matched with e_i) to pair with another incorrect power event, resulting in a series of mismatched pairs.

Figure 11 shows the statistics of the appliance operating duration resolution error of the HeatProbe system. The x -axis represents the amount of operating duration resolution error between the detected and actual operating durations, while the y -axis represents the cumulative distribution function of the operating duration resolution error. From the

Scenarios	Appliance	True Power [W]	Estimated Power [W] (Error)
1	PC #1	86.09	79.25 (7.94%)
	PC #2	79.91	74.45 (6.83%)
	LCD Monitor #1	27.86	40.08 (43.88%)
	LCD Monitor #2	26.94	39.34 (46.04%)
	Lamp #1	80.03	110.42 (37.97%)
	Lamp #2	29.61	82.04 (177.05%)
	Printer	256.23	286.42 (11.78%)
	Toaster	629.09	616.10 (2.06%)
	Heater	411.76	353.11 (14.24%)
2	Notebook #1	32.26	35.08 (8.77%)
	Notebook #2	29.86	29.80 (0.20%)
	Lamp #1	80.23	82.20 (2.45%)
	Lamp #2	82.44	88.11 (6.89%)
	TV	121.14	121.54 (0.33%)
	Bug zapper	33.65	30.32 (9.91%)
	Shredder	76.86	75.75 (1.45%)
3	Water heater	433.98	417.55 (3.79%)
	Microwave oven	733.15	741.23 (1.10%)
	Toaster	540.63	631.93 (16.89%)
	TV	121.28	121.61 (0.27%)
	Oven	547.85	457.87 (16.42%)

Table 4. Appliance power accounting error for all scenarios. True (Estimated) column means the true (estimated) power.

usage durations inferred from the previous matching results, 80% of duration errors are less than 120 seconds, with an average error of 125.03 seconds.

Overall appliance power accounting accuracy

Table 4 presents the average appliance power accounting accuracy. The ground-truth (estimated) powers for each appliance are averaged and listed in the table. The overall average error between the estimated power and the ground-truth power is 19.8%. Most errors are less than 10%. However, there are a few cases in Scenario #1 where the error rates are above 30%, e.g., LCD Monitors #1 and #2, and Lamps #1 and #2. These large errors are due to mismatched on/off events. For example, the 177.05% inaccuracy of Lamp #2 in Scenario #1 is caused by mismatching Lamp #2 events to the power amount of a heater. The two appliances greatly differ in power consumption, resulting in a significant energy accounting error.

CONCLUSION

This study presents HeatProbe, a novel thermal-based energy meter system to disaggregate total energy consumption. Experiments conducted on three multi-user scenarios showed that the overall appliance power accounting accuracy reached an average of 80.2%. The system successfully monitored appliance usage durations with an average error of 125.03 seconds. Moreover, 80% of the measured usage durations were under 120 seconds.

We believe that the HeatProbe system offers an alternative and promising thermal-sensing approach to tackle the problem of disaggregating electricity usages. In the future,

we will improve disaggregation accuracy and explore the possibility of accounting per-user energy footprints by incorporating human identification schemes.

REFERENCES

- Behavior, Energy, and Climate Change Conference. <http://www.beccconference.org/>
- Gupta, S., Reynolds, M.S., and Patel, S.N. ElectriSense: Single-point Sensing using EMI for Electrical Event Detection and Classification in the Home. In *Proc. Ubicomp 2010*, ACM Press (2010), 139-148.
- Kim, Y., Schmid, T., Charbiwala, Z.M., and Srivastava, M.B. ViridiScope: Design and Implementation of a Fine Grained Power Monitoring System for Homes. In *Proc. Ubicomp 2009*, ACM Press (2009), 245-254.
- Rowe, A., Berges, M., and Rajkumar, R., Contactless Sensing of Appliance State Transitions. In *Proc. BuildSys 2010*, ACM Press (2010), 19-24.
- Cent-a-Meter, <http://www.centameter.com.au/>
- TED: The Energy Detective, <http://www.theenergydetective.com/>
- Jiang, X., Dawson-Haggerty, S., Dutta, P., and Culler, D. Design and Implementation of a High-fidelity AC Metering Network. In *Proc. IPSN 2009*, IEEE Computer Society (2009), 253-264.
- Hart, G. Nonintrusive appliance load monitoring. In *Proc. the IEEE*, 80, 12 (1992), 1870-1891.
- Norford, L. K. and Leeb, S. B. Non-intrusive electrical load monitoring in commercial buildings based on steady-state and transient load-detection algorithms. *Energy and Buildings*, 24, 1 (1996), 51-64.
- Puri, C., Olson, L., Pavlidis, I., Levine, J., and Starren, J., StressCam: non-contact measurement of users' emotional states through thermal imaging. *Ext. Abstracts CHI 2005*, ACM Press (2005), 1725-1728.
- Yun, C., Shastri, D., Pavlidis, I., and Deng, Z., O' game, can you feel my frustration?: improving user's gaming experience via stresscam. In *Proc. CHI 2009*, ACM Press (2009), 2195-2204.
- Larson, E.C., Cohn, G., Gupta, S., Ren, X., Harrison, B., Fox, D. and Patel, S.N. HeatWave: Exploring the Feasibility of Thermal Imaging for Surface User Interaction, In *Proc. CHI 2011*, ACM Press (2011), 2565-2574.
- Bayati, M., Shah, D., Sharma, M. Maximum Weight Matching via Max-Product Belief Propagation, In *Proc. ISIT 2005*, ARXIV (2005), 1763-1767
- Bradski, G., Kaehler, A., *Learning OpenCV: Computer Vision with the OpenCV Library*, O'Reilly Media (2008).
- Nilsson J. W., and Riedel S. *Electric Circuits, 8th Edition*, Prentice Hall (2007).
- Patel, S.N., Gupta, S., Reynolds, M.S. The design and evaluation of an end-user-deployable, whole house, contactless power consumption sensor. In *Proc. CHI 2010*. ACM Press (2010), 2471-2480.
- Stern, P.C. Information, incentives, and proenvironmental consumer behavior. *Journal of Consumer Policy*, 22, 4 (1999), 461-478.