Real-Time Interactive Visualization of Disaggregated Energy Usage

Eric Mullen
University of Washington
emullen@cs.washington.edu

Md Tanvir Islam Aumi
University of Washington
tanvir@cs.washington.edu

Will Scott
University of Washington
wrs@cs.washington.edu

ABSTRACT

In this work, we present a starting step towards building a platform for real-time interactive visualization of disaggregated energy usage. To enable real-time communication between a server and client, our system streams data from the server using web sockets. In addition to displaying real-time appliance level feedback, this system also provides aggregated per appliance measurement. Furthermore, it allows users to explore details and see both current and more-general per-day usage trends in their devices. We believe that our system will eventually enable HCI researchers to perform real-world deployments across multiple homes and study the effects of visualization on energy disaggregation feedback.

Author Keywords

Real-time visualization; interactive visualization; energy disaggregation.

ACM Classification Keywords

H.5.m. Information interfaces and presentation (e.g., HCI): Miscellaneous.

INTRODUCTION

Energy conservation and eco-feedback research continues to be an important focus in the HCI communities. Given that 28% of U.S. energy consumption is directly contributed by household activities [2], the home is a natural place to study. As a result, sensing disaggregated electricity usage in the home (*i.e.*, at the individual source level) has emerged as one particularly promising area. There are many promising techniques that allow the users to get real-time appliance level energy usage data for their home [1,5]. Consequently, there is an increasing need of visualizing this data in a meaningful way.

In this work, we focus on creating a simple and user-friendly visualization of real-time appliance level energy consumption data. The goal is to create a visualization that will help the end-users to reduce their electricity consumption. For example, if a user sees that her furnace is costing her a lot on her electricity bill, she can just switch to an energy-efficient one, which will pay for itself in a few years. This work will not only help the user save money, but will also help them reduce their energy footprint on the planet.

To achieve this goal, we developed a system that shows real-time energy disaggregation feedback to end-users. Unlike previous approaches which only focus on providing historical insight into previously collected static datasets [7], our system aims to show feedback in real-time. Therefore at any point of time, users can see their power consumption change based on appliances turning on or off. To the best of our knowledge, this is the first attempt to visualize disaggregated energy usage feedback in real-time.

Our system allows users both to explore individual data and see how individual devices contribute to their total energy bill, and also to get summary statistics about devices, like the estimated yearly cost of that device. These summary statistics show both aggregate estimates, and current data on the current cost-of-ownership for those devices. We believe that these items in particular can help to create awareness among general consumers about the implications of their energy usage. Users are especially encouraged to observe the temporal variation of energy consumption for each device with simple mouse-over of the runtime graph. Finally, our system uses easily perceivable visualization techniques to indicate the device state. Therefore at any time, users are also able to see the exact devices on/off in their household. We believe that this information could also provide us new insights about activity recognition based on appliance usage.

Studies have shown that a real-time disaggregated energy feedback system can reduce energy consumption by 20% [2]. However, there are certain design aspects to consider. To evaluate those design aspects in large scale, we believe that there is a great need for a simple, robust, and easily deployable energy visualization system for residential houses. While this is a much more complicated design problem, we foresee our work as the first step towards it.

RELATED WORK

The concept of visualizing real-time disaggregated energy usage is not very old. However, there are some interesting visualization techniques that approached this idea. For example, [6] is a very good example of visualizing disaggregated energy usage in United States. It enables the user to explore how much each appliance uses in Watts and Dollars. It also shows how much use a kilowatt-hour yields per appliance. However, this visualization is for awareness purposes and does not give any real-time feedback of per home energy usage. Therefore its not possible to figure out which appliances are on or off and how much power each appliance is consuming in the moment.

[7] is a great example of kitchen energy monitoring where users can get feedback of energy usage of kitchen appliances. While this idea is easily extensible for all the appliances, it focuses on a static dataset and hence unusable for real-time feedback. It must be mentioned that providing real-time feedback is very important as research [2,3] has shown that it saves up to 20% of total energy consumption. Finally, [8] explored some interesting visualizations for appliance level energy. But this visualization is also static and not applicable in our domain.

In summary, showing real-time appliance level energy usage feedback can help us to save our energy footprint on earth [3, 4] and currently there is no visualization platform that provides that. We believe that our system is the first step towards creating such a platform. Not only this will help end-users to reduce their personal energy consumption, but it will also help HCI researchers to perform large-scale deployment studying energy consumption behavior.

IMPLEMENTATION DETAILS

Dataset and Preprocessing

We used a popular Reference Energy Disaggregation Dataset (REDD) for this project. This dataset gives us several weeks of power data for six different homes. The data itself and the hardware used to collect it are described in [4].

Briefly, for each house, this dataset includes real power information for all of the channels (appliances). The first two channels always describe the power usage by two main feeds coming into the house. Therefore to get the actual power consumption, these values are aggregated into a single 'whole-house power consumption' value understandable by the user. This is due to the fact that power is delivered to homes in three wires: one at 120V, one at 0V (neutral), and one at -120V. One channel of main data is for the power through the 120V line, the other channel is for the power through the -120V line. All of the other channels show the power consumption for each appliance. So their added value is the total power consumption. Before building the visualization, we used our server to produce summarizations of power usage and explore the data set – which was too large to easily understand by hand. The server results showed that the dataset does adhere to these conventions.

Server Side Implementation

As our system demands real-time streaming, we used web sockets to keep a persistent connection open between the server and the client. The server side code was written using python, based on the tornado web framework. We implement a pub-sub based protocol over the websocket connection where the client can ask for either the summarization of a specific time range for a device, or can ask to subscribe to real-time updates for that device. This is a reasonably generic protocol, which we believe helps to define what is needed in a home monitoring data collection system in order to create effective visualizations. One important note is that our initial approach of sending down the full data of the interesting time period was abandoned because there was simply too much data. Having appliance usage data at one second resolution for an entire day overwhelmed the client, both in memory pressure, and in its ability to create canvases that could responsively interact to user input.

It should be noted that the dataset we used were collected in different time resolutions. For each of the appliances, the resolution was every three seconds. For the main two lines, data was collected and recorded every second. That is one of the reasons we used web sockets. Instead of continuous polling, a web socket interface only pushes the data to the client when it is needed, and thus increases the efficiency of the whole data transfer process.

Client Side Implementation

The client side implementation is a webpage made with HTML, CSS, and JavaScript. Rickshaw, a time-series graphing library extending D3 was used for the primary data graphs. One unfortunate consequence of rickshaw is that the generated graph must be fully re-rendered for each change in data. One improvement that would help with the library would be allow it to instead react to the change in its data series and only update the affected paths.

Since the data we used was a collection from earlier, due to the fact that we do not currently have access to a house live-streaming its appliance data, we had to translate times between now and the data set. The client uses a date known to be in the data set, but uses the current time, to realistically approximate the live stream of data at the current time of day. In a real system, the client would simply use the current time, rather than translating it back to a known previous day.

To calculate costs, the client uses keeps track of total wattage reported by the server, and uses the standard Seattle area pricing to generate prices. These prices were in-line with typically experienced prices, which again helps to confirm the correctness of the data processing pipeline.

Icons are courtesy of the Noun Project. There are numerous data labels that don't lend themselves incredibly easily to icons, but the extensive library of icons at the Noun Project was useful to help liven up the page.

RESULTS

We built a real-time visualization of disaggregated energy data. We tackled several key issues with the state of the art, including allowing real-time streaming of the data, as well as presenting users an intuitive and obvious mapping explaining which of their appliances are currently powered on. In addition, we built a system that lets users compare their energy usage to the previous day, by overlaying the current day's energy usage with the previous day's. We believe that these three pieces, along with an attractive and simple interface will enable users to easily and effectively obtain a high level of knowledge about their personal energy usage. Simply being able to view energy use in real time will give users a sense of immediacy, and tie the use of energy to their innate conception that this is something happening in the present. Letting them easily and independently see what is on and off will help them debug any instances where they believe they have powered off an appliance, where they have not. Letting users compare the usage of the current day to the previous day will help them compete against themselves to lower their energy usage.

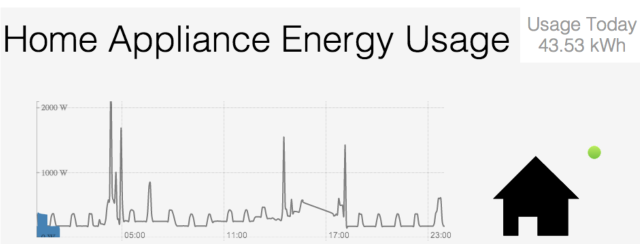

The primary view of our dashboard that a user is greeted with upon visiting the website has several important features:

It shows the total power used for that day in the upper right corner, for a single-number measure of energy usage.

The main graph shows the current and previous days energy usage, allowing a viewer to see both their energy usage over time, and what to expect in power consumption over the rest of the day.

Finally, the green dot indicates that the device is in use, which in the case of the whole-house means that the sensor is functioning correctly.

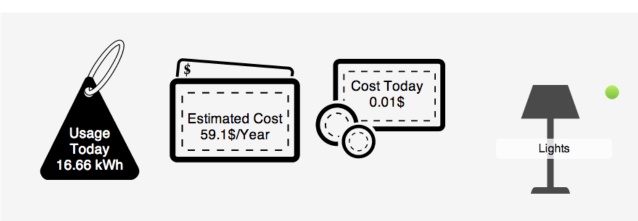

When drilling down into a specific device by clicking on that line in the dashboard, additional summary statistics are shown for that device. In this example, the house lighting is shown to cost around $60/year, and had used 16 kWh in the viewed day. We expect that these summary statistics can also be expanded, to include emissions, time to pay for a replacement, or other appropriate information.

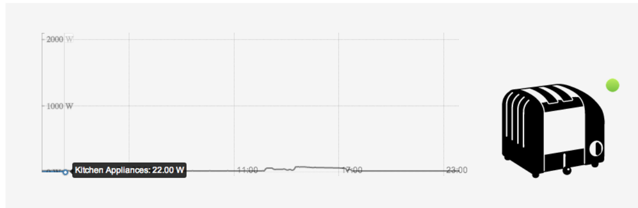

One of the coolest ways to explore the data is to play with the main graph itself. The graph has an interactive tooltip to allow the user to see specific energy usage values as well as how that usage has changed over time. The interesting point here is that many devices do not behave in the way that a user might initially expect. For example, while it may make sense that the the periodic cooling cycle of a refrigerator is correlated with energy draw, that is likely not something users necessarily have connected directly to their houses energy draw. Likewise, the relative magnitudes of something like a hairdryer may catch many users off guard, and being able to see how that compares to lighting or other devices that would not immediately seem to use significantly less energy may be an effective way to convince users to replace inefficient devices.

DISCUSSION AND FUTURE WORK

Our initial idea included a personalized feedback for each user. For example, the system could suggest a user to change her dryer and buy a energy efficient one, which will pay for itself in a few years. In fact, given dataset for all the houses in a neighborhood, it can also give such recommendation compared to one’s neighbor. We hope to explore this more on future.

Our initial prototype was actually designed to show feedback for every second. While this was a good technique to show actual real-time feedback, it was not useful for analyzing consumption behavior. We got some great feedback during the poster session regarding this issue, and decided to show summary in aggregated timescale. In future, we plan to incorporate both of these timescales into our design using simple drop down menu. This will enable the user to choose the preferred one based on her requirement. In fact, we also plan to perform an aggregation per day/week/month once we have access to a larger dataset.

Of course, in the future we would love to hook this into a physical tool that measures energy usage. This is near on the horizon, as a physical device to do just this is being released by Belkin later this year. Unfortunately we have preempted the release by enough that we weren't able to develop the visualization against an actual device. We are confident that there will be a low implementation cost to switch from the dummy streaming python server we have, to a device collecting data in your home.

There are additional challenges to appropriately visualizing home energy challenges – both in the form of creating action on the results provided by the view, and more easily communicating energy usage to the user. One feature that should be more clearly displayed is the current energy usage of each device. It seems like it could be natural for that value to appear as a large callout at the current time on the graphs, which also could highlight whether the appliance is on or off. However, the worry is that for the current energy usage to be readable, you would end up obscuring the most interesting part of the graph, and negate some of the value of seeing what the likely daily pattern is for the device.

A second addition would be to include not simply the previous day, but either multiple previous traces across the previous week, or an averaged trendline for devices. All of these approaches have advantages and disadvantages. Our current approach requires less data, allowing for cheaper tracking hardware and a simple to understand picture. It also, however, results in a rather jagged or potentially uncharacteristic view of energy usage. Using an averaged graph from more days has the advantage of smoothing out the data and being somewhat more representative about how energy usage typically is spent. It does not however communicate potential deviance, and may be bad at showing actual impact of usage if a device is not used at the same time across many days. Still, there is certainly room to improve in how we communicate to the user what the standard impact of a certain device is on their home's total energy bill.

It would likewise be cool if the users could select and visualize a section of energy, to try to get a feel for how much a particular use of something costs. There are many good spikes in our data when someone turns on a hair dryer or something like it, and it would be interesting if you could select that spike, and see statistics about it. In general people have a very bad impression of what uses power, and how much power each appliance uses. It would be cool to give users a hands on tool to take an amount of power they can conceptualize, e.g. the amount used to dry their hair in the morning, and give them the cost of what they just used right then.

Finally, it would be awesome if we could make this into a social application. It is cool that users could compete with themselves on a previous day, and try to better their energy use, but it would be even cooler if they could compete with their friends for dominance in energy saving. The best way to do this is still nothing well understood, but users tend to compete harder in a social environment, rather than alone against themselves. Ultimately the goal is to turn energy saving into a fun and competitive game, which inspires creativity and innovation in people's lifestyles.

CONCLUSION

We presented a first step towards building a real-time disaggregated energy consumption feedback platform. Our solution utilizes web socket to stream server data in real-time, which reduces network congestion significantly. In addition, it also leverages Rickshaw framework of D3 to create interactive graphs.

We consider our system as an initial approach towards exploring a decade-old energy-monitoring question: what HCI design aspects are more useful from an energy monitoring point of view? Therefore during the implementation, we tried to develop our system as a robust and flexible platform for further additions. Although current system is minimalistic in terms of feature, we believe that this is a great platform for conducting real world deployments across multiple homes. We hope to iterate and refine this system more in future with feedbacks from those deployments.

REFERENCES

1. Award, B.P., Gupta, S., Reynolds, M.S., Patel, S.N., and Science, C. ElectriSense : Single-Point Sensing Using EMI for Electrical Event Detection and Classification in the Home Best Paper Award. (2010).

2. Bin, S. and Dowlatabadi, H. Consumer lifestyle approach to US energy use and the related CO2 emissions. *Energy Policy 33*, 2 (2005), 197–208.

3. Dietz, T., Gardner, G.T., Gilligan, J., Stern, P.C., and Vandenbergh, M.P. Household actions can provide a behavioral wedge to rapidly reduce US carbon emissions. *Proceedings of the National Academy of Sciences of the United States of America 106*, 44 (2009), 18452–6.

4. Kolter, J.Z. and Johnson, M.J. REDD : A Public Data Set for Energy Disaggregation Research. 1, 1–6.

5. Patel, S.N., Robertson, T., Kientz, J.A., Reynolds, M.S., and Abowd, G.D. At the Flick of a Switch : Detecting and Classifying Unique Electrical Events on the Residential Power Line (Nominated for the Best Paper Award). (2007), 271–288.

6. http://visualization.geblogs.com/visualization/appliances/

7. http://fathom.info/kitchen

8. http://www.energycircle.com/blog/2010/02/12/powerhouse-dynamics-emonitor-progress-report-screen-shots-and-early-data