Title
Connected thermostats for low income households: Insights from user testing

Authors
Alston-Stepnitz, Eli
Sanguinetti, Angela
Peffer, Theres
et al.

Publication Date
2020

Peer reviewed

Connected Thermostats for Low Income Households: Insights from User Testing

Eli Alston-Stepnitz, UC Davis Energy & Efficiency Institute, Consumer Energy Interfaces Lab
Angela Sanguinetti, UC Davis Consumer Energy Interfaces Lab
Sarah Outcault, UC Davis Energy & Efficiency Institute
Therese Peffer, California Institute for Energy and Environment

ABSTRACT

HVAC energy use is typically the largest energy load in low-income households, which have fewer plug and appliance loads and suffer from substandard HVAC systems and poor building envelope. Connected thermostats, with features like remote control and machine learning, can be a cost-effective strategy to reduce HVAC energy use, but these devices are prohibitively expensive for low-income households. This research was part of a larger project to develop and test an affordable connected thermostat designed for low-income households. This paper presents results from usability testing of the thermostat mobile app and two iterations of the thermostat hardware with a total of 27 households from two low-income communities in Sacramento, California. We present findings regarding user preferences for using the app versus the hardware for different functionalities and the value of remote control via Bluetooth rather than Internet. Usability implications of trade-offs between functionality, cost, and design are discussed.

Introduction

Low-income households spend, on average three times more of their income on energy bills-- more than any other income group (Berry et al., 2018; Drehobl and Ross, 2016). Despite advancements in technology, the energy burden for low-income households has remained high, particularly among minority and rural households, and especially those with children and elderly residents (Lapsa et al. 2020). Heating, ventilation, air conditioning (HVAC) is typically the largest energy load in low-income households, accounting for 55% of total source energy and 30% of electrical energy use (Narayanamurthy & Zhao 2019). Programmable thermostats (PTs) offer some energy savings (Perry et al. 2011), but pale in comparison to more significant savings offered by more advanced, connected thermostats (Huchuk et al. 2018), which deliver accessible control through web, mobile, or voice platforms, and additional features for energy savings and load management (e.g., adaptation to occupant schedules). Thus, low-income households could potentially benefit from connected thermostats, but most available products are expensive and require access to WiFi (Ford et al. 2017; Pritoni et al. 2017), which may be less available in low-income households (Lapsa et al. 2020).

The energy impacts of programmable as well as more advanced thermostats are dependent on user behavior (Sanguinetti et al. 2017, 2018), which is heavily influenced by product usability (Nielsen 1993; Peffer et al. 2012; Perry et al. 2011). For example, actual outcomes of PTs, which require significant user interaction, tend to fall short of their energy savings potential because users find PTs difficult to understand and program (Gunshinan 2007; Peffer et al. 2011; Pritoni et al. 2015). Connected thermostats, while not as dependent upon user behavior as PTs if they include machine learning functions, still require user interaction with a mobile app and for certain functions, the device itself. Considering the rapid growth of the smart thermostat market[1] (Mordor Intelligence 2020), it is likely that manufacturers have conducted usability research, but as Meier and colleagues have noted (2010), insights from such studies are not available in the academic literature.

Looking to available industry research, Herter and colleagues (2014) conducted usability testing with a variety of commercially available smart thermostats. They found that several features contributed to user preference and ease of use, including large screen and menu text size, colored screens and overall aesthetics, and few required inputs to change modes. They concluded that utility programs involving thermostats should require threshold usability scores, conduct usability testing for all models being considered for programs, and provide extra training for renters and the elderly who take longer to complete tasks.

This paper presents insights from usability testing of a thermostat system called LYDIA, developed for the California Energy Commission project titled "Intelligent HVAC Controls for Low-Income Housing." While this research does not describe the product development phase, the following section provides a brief overview.

Background

Prior research has documented the energy-saving potential of advanced, connected thermostats. However, several factors inhibit their suitability for low-income housing environments, including high costs, wiring limitations, and limited access to WiFi. To meet the needs of low-income households and low-income multi-family building owners, the project team aimed to develop a thermostat that cost less than $75, did not require WiFi, and could be powered for up to one year on AA batteries, thus eliminating the need for a common wire or C-wire that many older HVAC systems do not have.

The thermostat was to be designed to control basic HVAC operations for up to 3-stage heat/2-stage cool, with humidity sensing, options based on energy costs, and fresh air control using ASHRAE 62.2. In addition, the original intent was to embed efficiency analytics based on consumers' revealed preferences, runtime data, indoor set points, indoor temperature, and stage operation. Another key objective was to ensure that the thermostat would allow users to achieve their comfort and energy- (or money) saving preferences. The intention was to design the hardware and software to be usable even with limited English and/or computer skills. Another objective was to incorporate features and functions that would deliver energy savings without the user having the explicit goal to conserve energy.

The outcome was two iterations of the hardware and a mobile application developed to work with the second iteration of the hardware. The second iteration of hardware, LYDIA, included some improvements based on usability testing (described in this study) and some to better meet the aforementioned objectives (e.g., a switch from LED light indicators to "e-paper" for the interface in order to save power). Features intended to save energy included on/off options for the whole HVAC system, timer, away mode, scheduler, and comfort band. The on/off, timer, away, and schedule features were intended to promote energy savings by keeping the system from running unnecessarily. The comfort band concept was to create a graphic display of temperature instead of a single digital display in order to show the current indoor

[1] The smart thermostat market was valued at USD 849.14 million, in 2019, and is expected to reach a value of USD 2858.63 million by 2025, at a CAGR of 23.1% during the forecast period (2020- 2025).

temperature within a range of comfortable temperatures (such as between the heating and cooling setpoints). The design intention was to promote the idea that there is no single temperature (i.e., setpoint) associated with comfort. While the original intent was to include machine learning, those functionalities were not ultimately developed due to data constraints. However, the unit was a connected device, able to use Bluetooth to communicate with a mobile application, thus providing the capability of machine learning in the cloud.

Methodology

Usability testing was conducted with the three LYDIA interfaces: two iterations of the thermostat hardware and one version of the companion mobile application. Testing was conducted with residents of two low-income multifamily housing communities[2] with central heating and cooling. Both communities are in the Sacramento, CA, area and managed by the same property developer, who was a partner in the research project.

The first round of testing focused on the first hardware iteration and was conducted at both communities, in August, 2018, with a total of 16 participating households. The second round tested the second hardware iteration and the mobile app, and was conducted at only one community, in December, 2019, with 11 households. Staff at each community assisted in recruiting participants by distributing flyers and scheduling resident appointments with researchers.

Testing took place in a quiet common area in each community. Two researchers were present for all test sessions: one directed the testing while the other took notes and assisted with recording equipment. Residents were scheduled to come in turns and invited to come alone or bring multiple household members (most came alone, with the exception of those needed translation support). First round testing sessions ranged from approximately 10-30 minutes in length and second round sessions ranged from 30-60 minutes. Each participating household received a gift card: $25 Target or Walmart for the first round (their choice), and $30 Target for the second round.

Participants included a range of ages and household compositions (e.g., single, couples, with and without young children, teenagers, and other arrangements of multiple generation households). Most participants spoke English, but participants who did not, or had limited fluency, brought a family member to translate or housing staff assisted. Several participants had mobility impairments.

Testing sessions proceeded as follows. Researchers asked participants to complete tasks with the thermostat system in order to assess the usability of each available feature and functionality on the hardware and/or app being tested. The first round of testing included five tasks to be performed using the first iteration of the hardware. The second round included the same five tasks to be performed with the second iteration of the hardware, and another eight tasks to be performed with the mobile application. The hardware and app were tested individually, not in conjunction, so interoperability issues were not addressed. In order to minimize potential bias from order effects in the second round (i.e., exposure to the app first versus the hardware first), half of the participants performed tasks on the hardware first and the other half on the app first.

The authors of this paper devised the testing protocols and the product development team reviewed them to ensure all functionality of the hardware and app were tested. A moderated

[2] Units subsidized by the Low Income Housing Tax Credit (LIHTC) program

“think aloud” protocol was employed for task completion activities. Researchers provided the initial instruction to complete each task and encouraged participants to speak aloud their thoughts and describe what they were doing as they attempted to complete the task. The researcher delivered minimal prompts and feedback as needed during task completion. Participant interactions with the device and their verbal comments were recorded. Video recordings were aimed at the device and participants’ hands; faces were not recorded.

Interspersed with the task completion activities, researchers posed questions to participants regarding how they would use certain features if the system were installed in their home, and asked for their subjective appraisals of various aspects of the system. The latter was also accomplished via post-testing interview questions in both rounds, and a usability survey was added to the protocol for the second round. The survey consisted of a scale called the System Usability Scale (SUS), which has ten questions, each with five response options, which aimed to quantitatively assess general usability of the LYDIA thermostat system (hardware and app). In their 2015 study, Herter and Okuneva suggested the SUS was a faster, cheaper alternative to extensive lab-based testing; their findings revealed a 91% correlation between the SUS and a more extensive usability scoring system.

Gazepoint eye-tracking equipment was used for usability testing with the mobile application to yield additional data regarding user experience. This equipment tracked participants’ eye movement to map their gaze on the screen displaying the mobile app. The mobile app was displayed on a computer monitor. This method enabled analysis of locations on the app screen that received more or less attention during each task.

Results and Recommendations

Results of both rounds of usability testing are summarized below along with interpretation and recommendations for future iterations of this product or other connected thermostats. Results are organized by usability tasks, followed by a section on further findings from post-testing interview questions, and the SUS score and comparison.

Turning Thermostat On

Across all three platforms, most participants were able to recognize the power button when prompted to turn the thermostat on. With the first hardware, a few participants had trouble seeing the power button clearly; thus, a recommendation from the first round of usability testing was to increase the contrast of the power symbol on the button. The second hardware iteration addressed that concern and no such issues were observed in the second round of usability testing.

Another issue with the first hardware was that some participants could not tell the difference between on and off states due to the illuminated green light to indicate the current temperature, which was present in both the on state and off state. The second hardware addressed this issue as well in that no information is displayed in the off state. Although this made it easy to determine whether the device is on or off, it precludes the ability for users to determine the temperature without turning on the system. Presumably for similar reasons as the first hardware (information displayed in off state), users occasionally had difficulty recognizing if the app was On or Off; however, the red/green color of the power icon in combination with the explicit “Turn Off/On” label beneath it did help users identify the state. Another cue could be added to make the On/Off state more obvious in the app, such as shading the screens except the power button when Off.

Reading Current Temperature

Across the three platforms, users easily identified the current temperature (Figure 1). The green light on the device in the first-round usability testing, which was present in the Off state, made it very clear to users. Most users correctly identified the current temperature on both the second device and app, however a couple participants did confuse one of the setpoints with the current temperature during testing (thinking the large number in the center might be a setpoint). The words heat/heating or cool/cooling directly beneath the indoor temperature may have contributed to this confusion. Potential solutions are placing the words "Heat to" or "Cool to" next to the corresponding setpoint or putting the setpoint in the phrase, e.g., "Heat to 72".

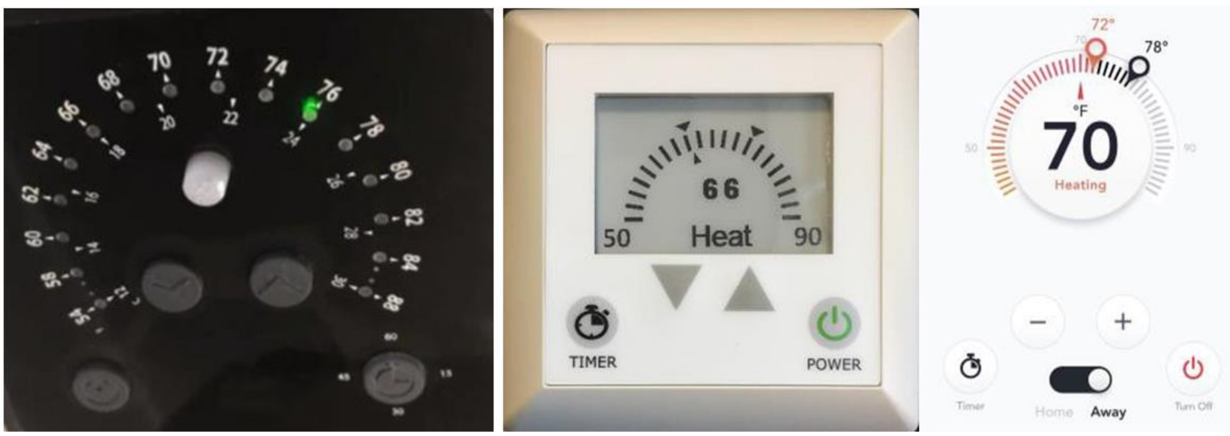

Figure 1. Current temperature (pictured left to right: first hardware; second hardware; mobile application).

Changing Temperature

Across the three platforms, users instantly understood how to increase or decrease temperature using the arrow or + and - buttons. In particular, users responded well to the immediate responsiveness of the first hardware and the mobile app to each button press, along with their numeric labels of each degree (first hardware) or setpoints (app); see Figure 2. However, on the second hardware the long lag between pressing an arrow and seeing the temperature change created a lot of confusion; the setpoints disappearing before they moved added to the confusion. Participants would often just keep pressing the arrows when there was no immediate response, in many cases resulting in a greater setpoint change than required to complete a task (e.g., turning to a lower setpoint than they might otherwise in order to turn on A/C); see Figure 3. They did not determine in advance the number of presses required for a desired change, which requires analyzing the degree scale (2 degrees between tick marks) and then counting or estimating the number of degrees between the desired setpoint and the current one. One participant remarked, “It needs to show you the temp moving as you click buttons.”

Figure 2. Labelled setpoints that change immediately upon user action on first hardware.

Figure 3. Setpoint change process on second hardware: setpoints disappear, then blinking box around new setpoint location, then reappearance of setpoints.

Moreover, although the arrow and +/- buttons were intuitive, most participants were confused by the presence of two setpoints (heating and cooling) on all three interfaces. Heat+Cool mode was the default setting on all interfaces tested and in this mode there are two setpoints: heating and cooling. Almost half of first-round participants and more than half of second-round participants did not successfully complete the task of getting the heat or A/C to come on without further prompts from the researcher. They often only moved the setpoints to the point where one was beyond the current temperature instead of both as required. Even after the researcher explained the heating and cooling setpoints, participants often became confused about them again in subsequent tasks (e.g., using the scheduler). Misperceptions included that the temperature would range between the setpoints (e.g., heat past the heating setpoint as far as the cooling setpoint or vice versa). When in Heat Only or Cool Only mode, participants were able to successfully complete the task of turning on A/C or heater. Participants universally preferred the single setpoint of Heat or Cool Only mode to the dual setpoints of Heat+Cool mode.

An interface that can respond to user actions immediately is crucial for usability. The temperature gauge should also be clearly labeled so that users can easily set precise setpoints. Another strong recommendation is to eliminate the Heat+Cool mode and dual setpoint approach, or at least do not make it the default setting. Instead, the defaults should guide the user to choose Heat Only or Cool Only mode or automated features could select for the user. This would eliminate confusion over the dual setpoints.

Another possible strategy to promote the idea of a comfort band is limit controls to setting a setpoint range rather than a single setpoint. This range could correspond to the hysteresis setting, the top or bottom of which would be the cooling or heating setpoint and the opposite end would be the point at which the A/C or heater kicks on (in Cool Only or Heat Only mode respectively). An eco-mode could increase this range. This would be educational, by promoting accurate perceptions that the temperature will fluctuate a little (which it always does anyway) and could thus help users become less attached to a specific single setpoint, which could ultimately pave the way for more conservative HVAC operations.

Setting Timer

The first hardware and the app allow for setting a timer for 15, 30, 45, or 60 minutes, whereas the second hardware only enables a 15-minute timer. Half of the first-round participants were able to complete the task involving the timer with the first hardware without any assistance. Some users had trouble seeing the button clearly. A few tried to turn the timer button rather than push it. It took a little experimentation or guidance to realize the button could be pressed multiple times to increase the timer duration (indicated by the LED orb around the button lighting up in quarterly increments with each press, and turning off with the fifth press). Some users misunderstood the function to be scheduling the system to turn on rather than off.

The timer button design was more salient in the second hardware and app design and labelled "Timer" (Figure 4). Again the lag on the second hardware confused users because they expected an immediate response. Only one participant accurately determined that the timer on the second hardware was for 15 minutes. Other guesses included 10 minutes (the icon on the digital display resembled a clock in the two-o-clock position).

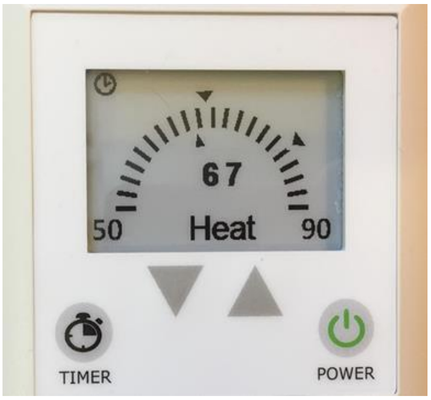

Figure 4. Timer icon displayed when the timer is set on second hardware.

On the app, like the first hardware, participants did not quickly realize they could set the timer for different intervals and they did not always notice the banner that appeared at the top of the app counting down the time. Some participants tried to increase the timer duration by pressing the + button or going into settings (the settings icon is near the timer banner). Similar to the first round, a few second-round participants misunderstood the function to be scheduling it to turn on rather than off.

Regarding the utility of the timer, about one-quarter of first-round participants said they would use it at night. Only one second-round participant gave that use case (perhaps this was due to the seasons, first round in summer and second in winter, if running the A/C at bedtime is more common than running the heater). Another use case offered was to turn it on while bathing children so it would be warm when they got out. To better promote understanding and use of the timer function, it should be named something more suggestive of the potential use cases (e.g., Heat/Cool Boost, Quick Cool/Heat, Cool/Heat Blast). Information about the duration of the timer should also be better integrated with the timer button (rather than appearing in a separate banner at the opposite end of the screen).

Turning Thermostat Off

Most participants in both testing rounds were easily able to turn the thermostat Off when prompted to turn off heating or cooling. Most liked this feature and thought they would use it. Several participants noted the utility of this feature for enabling children and elderly family members in their home to complete the task of turning Off heating or cooling.

Away Setting (App Only)

The Away setting was not intuitive to users because the setpoints did not change in a meaningful way when they first switched to Away and there were no prompts to change the setpoints or indications of the purpose of the feature (Figure 5). Also, the Away setting always defaulted to Heat+Cool mode even if the mode set for Home was Heat or Cool. Switching to Away should automatically create a default setback for the user and a tutorial wizard should explain this feature when the user first explores it. It should also default to the same mode as Home settings.

Several users accidentally increased or decreased their setpoints in the Away setting, thinking they were adjusting them in Home. It was also not clear to participants that they could not turn on the Schedule in Away mode (they had to switch back to the Home mode). Similar to the recommendation for the Off mode, having a cue across all screens to indicate when the Away setting is active would be helpful.

Some participants said they would use the Away feature, but many did not find it relevant. Most participants expressed that they or other household members are always home, which makes the Away feature less useful. Of those who considered changing settings when leaving, most preferred to just turn the thermostat off. One participant noted that Away would be more useful in the summer. These findings validate the decision to not include the Away feature on the hardware, and only provide it on the app (if that). Further usability testing after the aforementioned improvements to the Away setting would help determine if this is a useful feature to retain in the app.

Figure 5. Home and Away mode in the app.

Scheduler (App Only)

All participants were able to locate the scheduler, which allows users to set desired temperature ranges for different times of the day for each day of the week. Most were able to program the schedule, with two points of confusion: dual setpoints and absence of end times for events. The confusion around the dual setpoints described with the changing temperature task was compounded in the scheduler. Even participants who eventually grasped the meaning of the two setpoints on the temperature gauge in the main screen were perplexed about it again, perhaps because it is shown in different formats compared to the gauge (Figure 6).

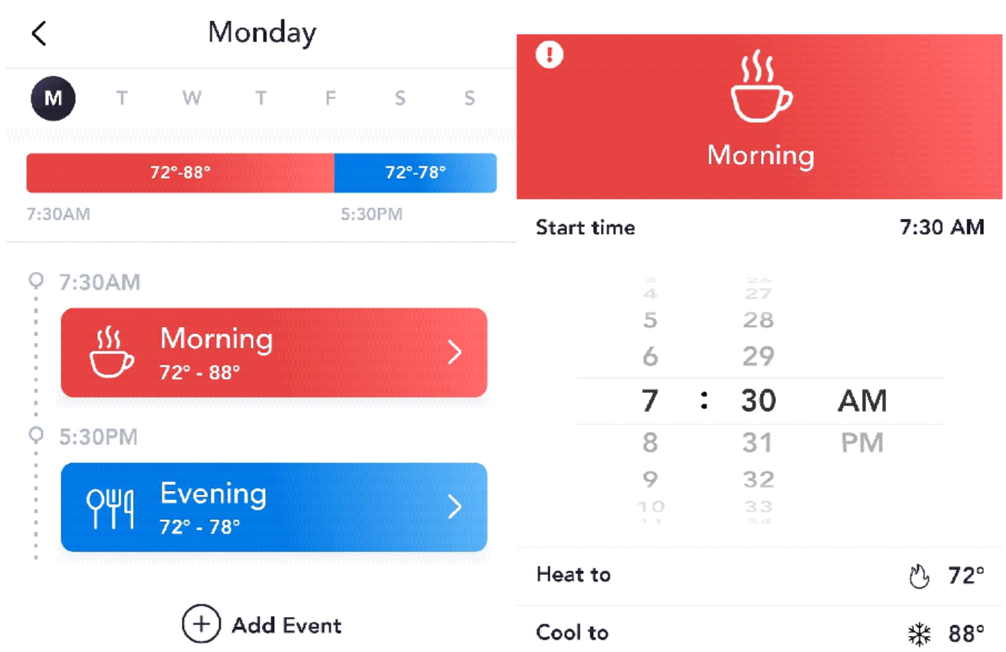

Figure 6. Reviewing and editing events in the schedule.

The "Heat to" and "Cool to" language in the editing events screen confused people; as one participant noted, "It's winter, so why would we be cooling?" Looking at event temperature ranges in the daily schedule (e.g., 72-82), participants often confused the cooling and heating setpoints or again misunderstood the setpoints to be a range within which the temperature would vary. Scheduling cooling and heating schedules separately, even when in Heat+Cool mode, would dramatically improve the usability of the scheduler.

The second issue, no end times for events, caused confusion because participants typically did not care to schedule an entire day, but instead had preferences for specific time periods. For example, several participants wanted to schedule the heater to come on in the morning and to a lesser extent in the evening but were less concerned with a concrete midday schedule. Once they put in start times for morning and evening events, and deleted daytime and nighttime events, they did not realize that the whole day up to evening event would be encompassed in the morning event and the whole night in the evening event. To fit user preferences and enable energy savings, the scheduler should allow gaps between events during which the thermostat would default to off rather than forcing setpoints for all times. This would also better promote energy conservation.

Usage Screen (App Only)

Most participants generally appreciated the usage screen (Figures 7-8), but many found it to be too much information and could not imagine specific uses for it. It reminded some of features on their energy bills. Percent difference compared to previous day, neighbors, or goal did not seem significantly or differentially motivating. Some expressed preference for comparing current month to same month last year. Most users said they would not be interested in seeing their usage compared with similar households or setting usage goals; it is worth noting that people tend to underreport the influence of social norms (Nolan et al. 2008). A few participants expressed that their energy costs are so low that this information is not motivating.

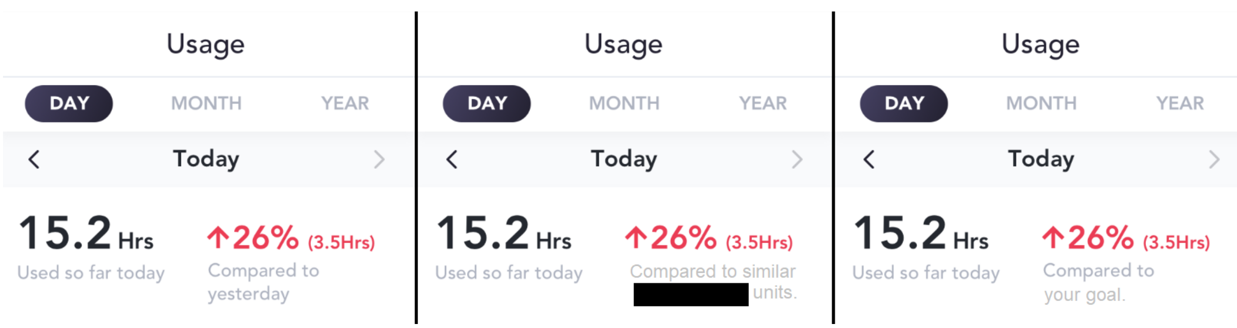

Figure 7. Usage screen stats (top of usage screen); participants shown three versions.

Providing strategically timed notifications during peak hours or perhaps just when users interact with the thermostat during peak hours or when they are using the scheduler would likely be more effective than a retrospective usage screen. Maintaining the usage screen seems worthwhile, but the information should be simplified, e.g., note heating and cooling usage separately in the stats at the top of the screen and exclude the Away and Off settings data from the bar chart. The monthly comparison view was not tested but might be more meaningful than daily since it can be used to gauge potential cost of bill compared to previous month. Comparing to same month of prior year, at the end of each month, as is common practice in energy feedback programs, could also be effective. A notification at the end of each month via the app to encourage users to check the usage screen could increase engagement.

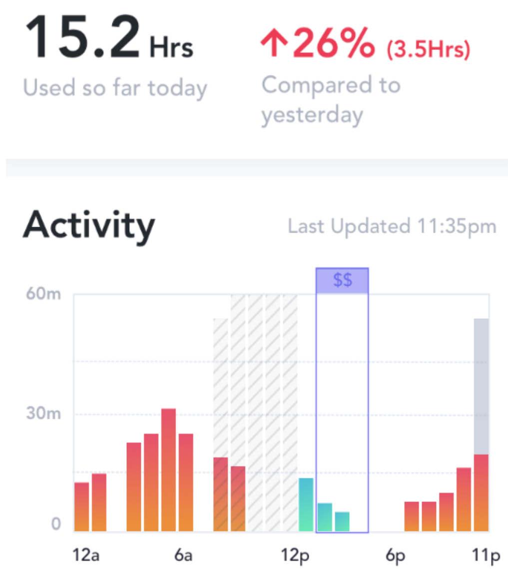

Figure 8. Usage screen graph (bottom of usage screen).

Post-Test Interview Questions

Most participants had a smartphone and liked the idea of having an app to control their thermostat. Only a few participants said they would not likely use the app. Some were disappointed that the app would not work outside the home, but most still appreciated it, particularly those with mobility impairments and elderly participants. With the exception of the second hardware, most participants preferred the LYDIA system to their current thermostat. Several participants noted that they would want or need training with the system before using it. Usability testing results should be used to inform training materials for thermostats.

System Usability Scale (SUS)

The average SUS score for LYDIA among second-round usability testing participants was 69. A SUS score of 68 is considered average. It is worth noting that most participants had trouble with the language in the SUS, either because English was not their first language, and/or the terminology was unfamiliar (e.g., terms like “cumbersome”). Furthermore, some participants expressed that their current thermostat was quite unsatisfactory and thus may have been judging LYDIA against a relatively low standard. These observations suggest caution when interpreting results and special consideration when using the SUS in future studies with populations in low-income housing or other underserved communities.

Conclusion

This research helps fill the gaps in academic literature on connected thermostat usability and thermostat usability among underserved populations. The findings highlight important features and considerations for usable, energy-conserving, and affordable connected thermostats, as well as other types of thermostats. Displaying heating and cooling setpoints at the same time was tremendously confusing. Further research would be needed to understand the degree to which this issue would persist with prolonged product use, other means of promoting the comfort band concept, and whether over time it could result in perceptual shifts among users to promote acceptance of a wider indoor temperature range.

The On/Off and Timer functions, which are simple yet uncommon features, were particularly popular. Participants found the ability to turn the system on or off desirable and were able to identify use cases for it that might result in energy savings compared to their current habits. The timer was generally appreciated and seems a promising feature for both user experience and energy savings given design improvements to clarify the function and suggest energy-saving use cases. On the other hand, participants expressed relatively less interest in more complex features: Away, Scheduler, and Usage. More research is needed to explore whether and how these features could be designed to be more meaningful for low-income households.

Results were consistent with Herter's (2014) findings that aesthetics (appearance, feel, and sound) are important. Participants were generally positive regarding aesthetics of the first generation hardware which included three dimensional buttons and colored LED light indicators. This overall positive response to the first hardware (tested without a companion app) demonstrates that an affordable connected thermostat device can be pared down to basic features and still be appealing. In contrast, participants were quite put off by the appearance of the second generation device and dissatisfied with its insufficient (tactile, visual, or auditory) feedback during interactions (compounding frustration with lagging response time). Given that the hardware interface was changed from LED lights in the first prototype to e-paper in the second iteration to save power, it seems that designing a product to run on batteries for applicability to buildings without a C-wire connection may not be worth the necessary sacrifices in interface design and resultant usability. Further research with alternate designs is needed to confirm this.

Almost all participants' valued the app as the strongest feature of the system—something that both increased convenience as well as removed physical ability barriers— though some felt limited by Bluetooth control and would prefer a WiFi enabled system that would allow them to control their thermostat while away from home. Overall, the findings suggest that connectivity without WiFi is still worthwhile for this population and likely others.

In sum, despite limitations of the LYDIA system, usability participants in low-income communities were overwhelmingly positive about the idea of a connected thermostat with a basic hardware interface and a companion app. The system's simple yet relatively novel features (easy On/Off and Timer) were well-received. Given the potential for connected thermostats to support enhanced comfort and energy and cost savings among low-income households, further efforts are needed to made these technologies more accessible to them and suited to their needs. To achieve these goals, usability considerations should be integral to product design, development, and evaluation.

References

Berry, C., C. Hronis, and M. Woodward. 2018. "Who's Energy Insecure? You Might be Surprised." 2018 ACEEE Summer Study on Energy Efficiency in Buildings. https://aceee.org/files/proceedings/2018/index.html#/paper/event-data/p393

California Budget Project. 2007. *A Generation of Widening Inequality: The State of Working California, 1979 To 2006.*

Churchwell, C., and M. Sullivan. 2014. *Findings from the Opower/Honeywell Smart Thermostat Field Assessment.* PG&E Emergent Technologies Program. Project Number: ET11PGE3074

Drehobl, A., and L. Ross. 2016. “Lifting the High Energy Burden in America’s Largest Cities: How Energy Efficiency Can Improve Low Income and Underserved Communities.” American Council for an Energy-Efficient Economy (ACEEE). https://aceee.org/researchreport/u1602

Ford, R., M. Pritoni, A. Sanguinetti, and B. Karlin. 2017. “Categories and Functionality of Smart Home Technology for Energy Management.” *Building and Environment*, 123, 543-554.

Gunshinan, J. 2007. “Energy Star Changes Approach to Programmable Thermostats.” *Home Energy Magazine*.

Huchuk, B., W. O'Brien, and S. Sanner. 2018. “A Longitudinal Study of Thermostat Behaviors Based on Climate, Seasonal, and Energy Price Considerations Using Connected Thermostat Data.” *Building and Environment,* 139, 199-210.

Herter, K., and Y. Okuneva. 2014. *SMUD’s Communicating Thermostat Usability Study*. El Dorado Hills, California: Herter Energy Research Solutions. Accessed from: www.HerterEnergy.com/pages/publications.html

Lapsa, M.V., M.A. Brown, A. Soni, and K. Southworth. 2020. *Low-Income Energy Affordability: Conclusions From A Literature Review*. (No. ORNL/TM-2019/1150). Oak Ridge National Laboratory (ORNL), Oak Ridge, TN (United States).

Meier, A., C. Aragon, B. Hurwitz, D. Mujumdar, T. Peffer, D. Perry, and M. Pritoni. 2010. *How People Actually Use Thermostats*. (No. LBNL-4977E). Lawrence Berkeley National Laboratory (LBNL), Berkeley, CA (United States).

Mordor Intelligence. (2020). *Smart Thermostat Market- Growth, Trends and Forecast (2020-2025)*.

Narayanamurthy, R., and P. Zhao. 2019. *Scalable Near Zero Energy Retrofits in Low-Income Multifamily Housing*. Energy Research and Development Division Final Project Report. California Energy Commission.

Nolan, J. M., P. W. Schultz, R. B. Cialdini, N. J. Goldstein, and V. Griskevicius. 2008. “Normative Social Influence is Underdetected.” *Personality and Social Psychology Bulletin*, 34(7), 913-923.

Nielsen, J. 1993. *Usability Engineering*. San Francisco: Morgan Kaufmann.

Peffer, T., M. Pritoni, A. Meier, C. Aragon, and D. Perry. 2011. "How People Use Thermostats in Homes: A review." *Building and Environment,* 46(12), 2529-2541.

Perry, D., M. Pritoni, C. Aragon, A. Meier, and T. Peffer. 2011. "Making Energy Savings Easier: Usability Metrics for Thermostats." *Journal of Usability Studies,* 6(4), 226-244.

Pritoni, M., R. Ford, B. Karlin, and A. Sanguinetti. 2017. "Home Energy Management (HEM) Database: A List with Coded Attributes of 308 Devices Commercially Available in the US." *Data in Brief,* 16.

Sanguinetti, A., B. Karlin, R. Ford, K. Salmon, and K. Dombrovski. 2018. "What's Energy Management Got to Do with it? Exploring the Role of Energy Management in the Smart Home Adoption Process." *Energy Efficiency,* 11(7), 1897-1911.

Sanguinetti, A., B. Karlin, and R. Ford. 2017. "Smart Home Consumers: Comparing Self-reported and Observed Attitudes." *Proceedings of Energy Efficiency in Domestic Appliances and Lighting (EEDAL).*