WHITE PAPER

Enhanced Auto-Schedule

Nest Labs, Inc.
November 2014

Introduction

In November 2014, Nest released Enhanced Auto-Schedule, an updated and re-designed version of schedule learning for the Nest Learning Thermostat. This flagship feature of the thermostat allows users to obtain the benefits of a heating and cooling schedule without the hassle of manually programming their thermostat. It adapts to customers' lives, making customers comfortable while saving energy at the same time.

In this paper, we describe how Enhanced Auto-Schedule allows customers to save even more energy by learning a schedule that better fits their routine. Specifically, we show in simulation that Enhanced Auto-Schedule saves users 23.1% in cooling and 19.5% in heating. Compared to our previous approach, this is an additional savings of 6.1% for cooling and 5.6% for heating. This estimate doesn't guarantee specific energy savings. Actual energy savings will depend on factors beyond the Nest Thermostat's control, such as furnace efficiency, home construction and weather, and customer interactions.

The simulations used in this paper were based on home data for North America, because a full season of heating data is not yet available for Nest Thermostats in European countries. However, we believe that savings in Europe may be similar to those in the US because Auto-Schedule is primarily based on learning from user inputs and target temperatures. In general, for countries with milder climates, savings numbers can be expected to be higher, as changes in target temperature will have a larger impact on HVAC runtime. For countries with more extreme climates, savings numbers may be lower. As Nest has done with other white papers, the Auto-Schedule white paper will be revised accordingly as more data becomes available for homes in European markets.

In the US, heating and cooling accounts for approximately 54% of residential energy [1]. In the countries where the Nest Thermostat is available in Europe, heating accounts for 44-65% of residential energy use (UK: 60%[2]; France: 61%[3]; Ireland: 44%[4]; Netherlands: 65%[5]; Belgium 54%[6]). Basic, non-programmable thermostats generally hold a "comfort temperature" for a majority of the time, wasting valuable energy during potential set-back periods, such as when a user is asleep, away, or otherwise willing to accept a more energy-efficient temperature.

Programmable thermostats are designed to save energy by allowing the user to enter a schedule. However, the user interface for entering schedules is often not intuitive or not flexible enough to meet user needs. This results in many users never actually entering a schedule. Furthermore, even when they put the schedule in, when behavior patterns change, the user resorts to a "hold" mode similar to an unprogrammed thermostat. In fact, many studies have documented that programmable thermostats are often operated in "hold" mode, and may produce no net savings over traditional non-programmable thermostats [8, 9, 10, 11, 12]. This failure of existing tools provides an opportunity for savings through the use of automatic schedule learning. Specifically, the EPA has estimated that having a properly programmed schedule with good setbacks can save up to 20% on heating and cooling bills [13].

Since its original release in October 2011, Nest Auto-Schedule has overcome these problems by learning from user inputs and other sensor data to develop a schedule that works for each home. It starts with an empty schedule so that it can learn the routines and preferences of the user, and continues to adapt to their ever-changing schedule. Our past study showed that Nest Auto-Schedule, together with Auto-Away, leads to 19.8% cooling and 16.2% heating savings during peak seasons, as compared to traditional thermostats holding constant at their personal 90th percentile temperature [14]. It is also the second most loved feature by Nest customers.

In our early trials prior to launching the first-generation Nest Learning Thermostat, we discovered two things: (1) Users need to be in control, and therefore Auto-Schedule should not choose a temperature the customer didn't select, and (2) Auto-Schedule should start with a blank schedule and adapt quickly at first. When we assumed that users wanted to save energy and helped them by turning down the temperature to a number they had never set, they were not happy with the outcome. Furthermore, when we started off with a set schedule as a default, users tended to over-adjust the temperature in such a way that they ended up with a schedule that wasted more energy than if the thermostat simply started with no schedule. Making users fight against temperature schedules they did not select or want caused not only irritation and discomfort but also thermostat usage that resulted in higher energy usage than before. By nature, people don't like being told what to do.

Throughout the past three years, Auto-Schedule has been steadily updated and improved, but Enhanced Auto-Schedule is our first major re-design based on customer feedback and analysis of data from the field across the US, Canada and the UK. Critical areas we have addressed with Enhanced Auto-Schedule include:

- Increasing efficiency while respecting user inputs and not compromising comfort, and
- Improving responsiveness of schedule learning to changes after the first few weeks, including changes of season and user routines.

How does Enhanced Auto-Schedule work?

The Enhanced Auto-Schedule algorithm is predicated on paying even more attention to user inputs: every interaction is treated as a way for the user to communicate with the device about his or her preferences for a particular temperature at a particular time and day of the week. In addition to considering active interactions, we also consider lack of interactions (indicating satisfaction with the current temperature), as well as the room temperature and whether the user is home or away. This provides a more holistic view of user preference than was considered previously.

Methods

We evaluated the effectiveness and improvements of Enhanced Auto-Schedule with simulation and in a field trial. We simulated the schedule learning for an aggregate set of devices using both the original and Enhanced Auto-Schedule algorithms. Our simulation data consisted of real, anonymized historical data of user behavior from the first four weeks of device installation. We considered devices that disabled Auto-Schedule soon after installation in order to simulate on temperature changes that are more indicative of actual user preference rather than responses to the learned schedule that they experienced. We simulated the schedule learning with logged user dial turns for the first four weeks following installation using both the original and Enhanced Auto-Schedule.

Second, in order to determine whether we were truly maintaining user comfort with this upgrade, we also ran a field trial in which we recruited users, told them that they were receiving a new version of Auto-Schedule, and gave half of them the original algorithm (the placebo group) and half of them the Enhanced algorithm (the experiment group). Our trial was conducted for six weeks during late summer when users were predominantly running air conditioning.

Results

Efficiency Improvement

Enhanced Auto-Schedule improved savings across the simulated devices. Savings are expressed differently for different devices. Some examples are shown below. In Figure 1, we show an example where the workday heating setback is longer under the new algorithm. This schedule resulted from paying careful attention to adjustments the user made.

In Figure 2, we show another example where the setback learned by Enhanced Auto-Schedule is both longer and deeper than it was previously. In addition, the comfort temperature is 1°F more efficient than the previously learned schedule.

Table 1 shows our quantitative findings across the full dataset, showing the savings of Auto-Schedule when compared to the thermostat holding the user's 90th percentile temperature during the first four weeks of device lifetime as described above. This is standard methodology that we used in our past report and is supported by independent analysis [14, 15]. Compared to this flat schedule, Enhanced Auto-Schedule saves users 23.1% in cooling and 19.5% in heating. Compared to the previous approach, this is an additional savings of 6.1% for cooling and 5.6% for heating. For completeness, we include the Auto-Schedule plus Auto-Away savings numbers during peak season as reported in our previous study [14] in the first column.

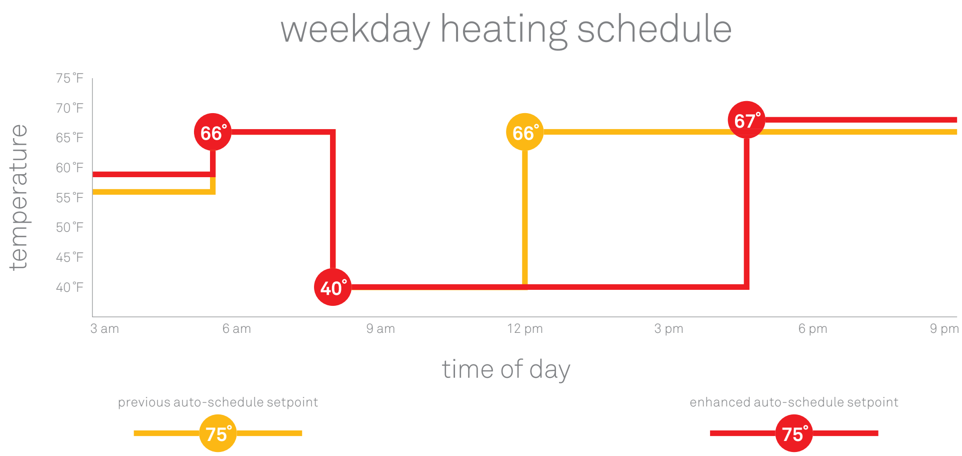

Figure 1: Enhanced Auto-Schedule achieves savings by choosing a longer setback while still respecting the user's dial turns and preference indications.

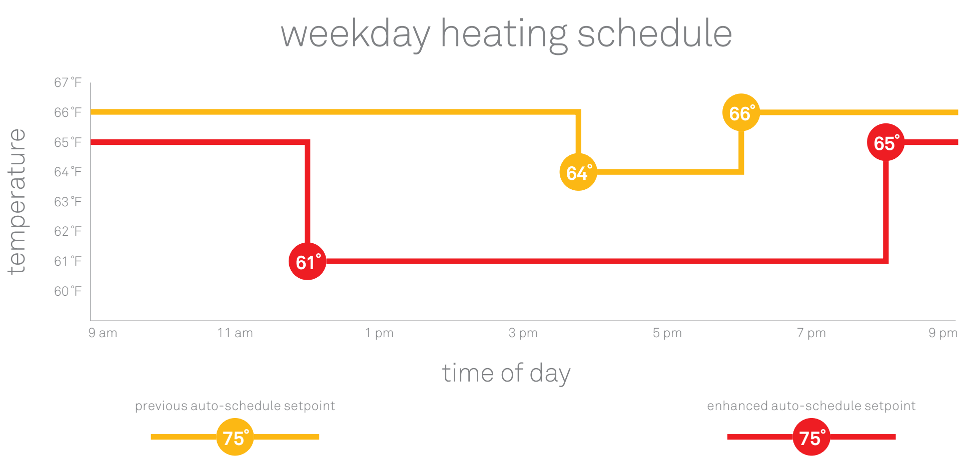

Figure 2: Enhanced Auto-Schedule learns a longer and deeper setback, reflecting more efficient settings that the user chose, saving substantial energy.

table 1

	auto-schedule with auto-away savings in peak season [14]	original auto-schedule savings over first four weeks	enhanced auto-schedule savings over first four weeks	overall improvement
cooling savings	19.8%	17.0%	23.1%	**6.1%**
heating savings	16.2%	13.9%	19.5%	**5.6%**

[savings from simulation comparing original to enhanced auto-schedule]

The savings are not in any way forced on every thermostat. If the user prefers more comfortable temperatures and adjusts the thermostat accordingly, schedule learning will not resist. However, if the user does make temperature changes to more efficient temperatures, Enhanced Auto-Schedule will pay more attention to these changes than previously.

Adaptiveness Improvement

Enhanced Auto-Schedule also adapts to users adjustments more closely when appropriate. By taking the entire history of the user's interactions into account in a holistic view, Enhanced Auto-Schedule does a better job of translating consistent and intentional temperature changes after the first few weeks into learned setpoints. In one example, a user installed Nest in mid December, preferring a morning (9 AM) temperature of 72°F. After a few weeks however, the user indicated with repeated dial turns that they instead preferred 69°F. While our previous algorithm would have missed this change, Enhanced Auto-Schedule captured this preference, changing the 9AM setpoint from 72°F to 69°F.

We measured the adaptability of the algorithm quantitatively in our field trial by comparing how often the placebo group or experiment group chose to adjust the temperature or to manually edit the schedule. Our findings show that users with Enhanced Auto-Schedule adjusted the temperature 8% less often compared to the control group. Furthermore, they edited their schedule 31% less often. These numbers indicate that Enhanced Auto-Schedule is better able to adapt to users' preferences than our previous approach. With a better adapted schedule, users are able to maintain their preferred temperatures with fewer dial turns and schedule edits.

table 2

metric	original auto-schedule	enhanced auto-schedule	improvement
number of temperature adjustments during the trial period	33.31	30.74	2.57
number of schedule edits during the trial period	1.69	1.16	0.53

[quantitative analysis of adaptability improvements of enhanced auto-schedule in field trial]

Overall Experience Improvement

Finally, we conducted a survey of our trial participants - both the placebo and experiment groups. The survey was sent by email and 83% responded. Detailed quantitative results can be seen in Table 3 below, and Table 4 contains some comments from users in the experiment group. 10% more participants in the experiment group preferred Enhanced Auto-Schedule than those in the placebo group who received the original Auto-Schedule. A small portion of both groups commented that they preferred the previous Auto-Schedule, at similar response rates.

table 3

which learning experience did you prefer?	original auto-schedule (placebo)	enhanced auto-schedule (experiment)
previous auto-schedule	14%	16%
no preference	50%	38%
enhanced auto-schedule	36%	46%

[survey results comparing user preference for new algorithm]

table 4

"I think this new version ran more smoothly, and I didn't have to remind nest that I would prefer a certain temperature at a certain time of day."

"Funny thing, first couple days I adjusted the temperature quite a bit (picky family) but since then I haven't touched it much, it's schedule worked very well considering how much we were adjusting it the first few days."

"The trial fine-tuned my schedule accurately and the temperatures were much more balanced throughout the day and week. I also noticed more Nest green leaves in the process."

"I felt the learning phase was very quick. Basically the next day I was noticing the set points. It felt easier than the first time."

[experiment group participant comments about enhanced auto-schedule]

Conclusion

In conclusion, Nest's new Enhanced Auto-Schedule more closely follows user behavior and intentions. This new schedule learning improves further upon the savings obtained by the learning approach pioneered by the Nest Learning Thermostat, the only approach proven in the field to capture savings in the large number of thermostats that are unprogrammed today. Specifically, we have shown that:

- Auto-Schedule has been saving users up to 19.8% cooling runtime and 16.2% heating runtime during peak seasons in the field compared to having no schedule.
- Our new Enhanced Auto-Schedule provides an additional 6.1% cooling and 5.6% heating savings on top of the existing savings provided by Auto-Schedule.
- Enhanced Auto-Schedule provides increased savings by understanding user preferences even better than before.

We look forward to continuing to incorporate user feedback in order to keep improving Auto-Schedule into the future.

The savings results reported in this white paper are based on simulations and our analysis. Actual savings will vary with a number of factors, including weather, your energy use, utility rates and plan. This is an estimate and not a guarantee of savings.

References

[1] US Department of Energy 2011 Buildings Energy Data Book, Chapter 2. http://buildingsdatabook.eren.doe.gov/ChapterIntro2.aspx

[2] Household energy bills explained. OFGEM; 2013. https://www.ofgem.gov.uk/ofgem-publications/64006/householdenergybillsexplainedudjuly2013web.pdf

[3] Agence de l'environnement et de la maîtrise de l'énergie. http://www.ademe.fr

[4] Eurostat. http://epp.eurostat.ec.europa.eu/

[5] Center for Energy and Environmental Sciences, IVEM. http://www.rug.nl/research/energy-environmental-sciences/?lang=en

[6] Energie Wallonie. http://energie.wallonie.be/fr/index.html?IDC=6018

[7] A Meier. How People Actually Use Thermostats. ACEEE Summer Study on Energy Efficiency in Buildings. Pacific Grove, Calif.: American Council for an Energy Efficient Economy. 2012. escholarship.org. http://escholarship.org/uc/item/3vd5q0cp.pdf

[8] T Peffer, M Pritoni, A Meier, C Aragon, D Perry. How People use Thermostats in Homes: A Review. Building and Environment, 2011 - Elsevier. http://eec.ucdavis.edu/files/How_people_use_thermostats_in_homes.pdf

[9] Haiad C, Peterson J, Reeves P, Hirsch J. Programmable thermostats installed into residential buildings: predicting energy savings using occupant behavior & simulation. Southern California Edison; 2004.

[10] Cross D, Judd D. Automatic setback thermostats: measure persistence and customer behavior. Chicago: IEPEC; 1997.

[11] Nevius M, Pigg S. Programmable thermostats that go berserk: taking a social perspective on space heating in Wisconsin. Proceedings of the 2000 ACEEE Summer Study on Energy Efficiency in Buildings; 2000:8.233e44.

[12] Shipworth M, Firth SK, Gentry MI, Wright AJ, Shipworth DT, Lomas KJ. Central heating thermostat settings and timing: building demographics. Building Research & Information 2010;38:50e69.

[13] Nest Learning Thermostat Efficiency Simulation: Update Using Data from First Three Months. Nest Labs. http://downloads.nest.com/efficiency_simulation_white_paper.pdf

[14] The Effect of Schedules on HVAC Runtime for Nest Learning Thermostat Users. Nest Labs. September 2013. http://www.efi.org/docs/The_Effect_of_Schedules.pdf

[15] A Data-Driven Framework for Comparing Residential Thermostat Energy Performance. Bryan Urban and Kurt Roth, Fraunhofer USA. July 2014. http://www.energystar.gov/sites/default/files/specs//Nest%20Supplementary%20Comments.pdf