ATherm®

AI-Driven Automatic Thermostat

Research Paper at Siemens PLM Software

by

Anish Bajaj and Pratyush Avi

June 2017

Contents

Introduction

This report includes the description of a project initiated and completed by Anish Bajaj and Pratyush Avi, two high-school students who sought to solve the problem of excessive energy consumption by modern air conditioners (AC's).

The product, A-Therm, was implemented by the duo at Siemens PLM Software, Global Business Park, Gurugram, India, during the summer of 2017.

It may be noted that this product was not always the vision we had in mind: we were only supposed to come up with an IoT Project during our internship at Siemens. For this, we started by surfing through various online forums like Kickstarter and Patreon to get an idea of the projects that had already been implemented. We wanted to create something unprecedented, something new.

A-Therm was subsequently pitched to venture capital investors at Siemens and received investment for further development, followed by its installation in the air-cooling system of the same branch in order to cut down on overall electricity consumption.

This report is the amalgamation of the ideation, implementation, working, usage, challenges, and outcomes related to the product.

The Problem

Modern air-conditioners work largely using thermostat systems. A thermostat is a component which senses the temperature of a system so that the it's temperature is maintained near a desired **set point**. The wide usage of this principle is proof that the system works up to one's expectations, but that does not mean that it is efficient altogether.

The efficiency of the aforementioned principle of thermostats is undercut by the indolence they provoke. If one sets the thermostat temperature to, for example, 18°C on a hot 35°C day, the air-conditioner will work perpetually until it reaches that set point. Although this may seem appropriate at first, one usually starts feeling too cold before the attainment of the set point because of the unabating working of the air-conditioner; in other words, the set point determines the temperature at which the air-conditioner would stop cooling, but not the temperature at which *the user would feel cool*. This eventually leads to the user's turning the machine on and off recurrently, and in turn results in poor efficiency due to the inability to keep a constant set point and forgetting about the heat/cold entirely.

But this is not the only drawback of the modern thermostat: because of the set point feature, a large amount of energy is utilized if and when the user chooses not to switch the air-conditioner off before the attainment of the set point – even if the user feels that the room air has sufficiently cooled.

Hypothesis: Thus, for an ideal thermostat system, the set point should be such that it cools the room as well as the user while consuming the minimum energy possible.

There are numerous ways to view this conundrum and there are numerous ways to solve it. We just had to look for the most elegant solution to this complex problem.

Our solution: A-Therm.

Introducing ATherm®

A factor that is usually ignored while calibrating air-conditioners is atmospheric humidity, which can make mildly hot temperatures unbearable, cold environments frigid, and moderate temperatures suffocating.

When calculated with respect to humidity, the temperature value so obtained is called the **Heat Index** (commonly known as **RealFeel® Temperature**, i.e., a measure of how hot it really feels when relative humidity is factored in with the actual air temperature.)

> "*The RealFeel Temperature is an index that describes what the temperature really feels like. It is a unique composite of the effects of temperature and humidity on the human body — everything that affects how warm or cold a person feels.*"

When an air-conditioner's thermostat is manipulated using the Real Feel Temperature rather than just the atmospheric value, it yields a much more accurate value for the required room temperature. Further, this value is made more precise by the usage of the **Thermal Comfort Temperature**.

"*Thermal comfort is the condition of the mind that expresses satisfaction with the thermal environment.*"

The so obtained temperature was termed the **Optimal Thermal Comfort Temperature** for the thermostat. At this configuration, a person present in the room feels neither too hot nor too cold. Apart from unburdening the user from fiddling with the air-conditioner remote to feel comfortable, this configuration can potentially reduce the electricity consumption and increase the longevity of the machine.

All of this combined information led to the creation of ATherm – Automatic Thermostat. It works on a very basic principle: receiving the values of the atmospheric temperature and relative humidity and subsequently calculating the optimal temperature.

Methodology

Immediately after coming across the idea, we began with the research for building the prototype: the preliminary research comprised of understanding the most basic factors that affect the environmental conditions of a room. From the information thus obtained we began hunting for parts for our model.

Materials Used

- Microprocessor
- Temperature Humidity Sensor
- USB Cable
- LCD Monitor
- WiFi Module
- Box unit (LEGO for the initial prototype)
- White case (final product)

Build

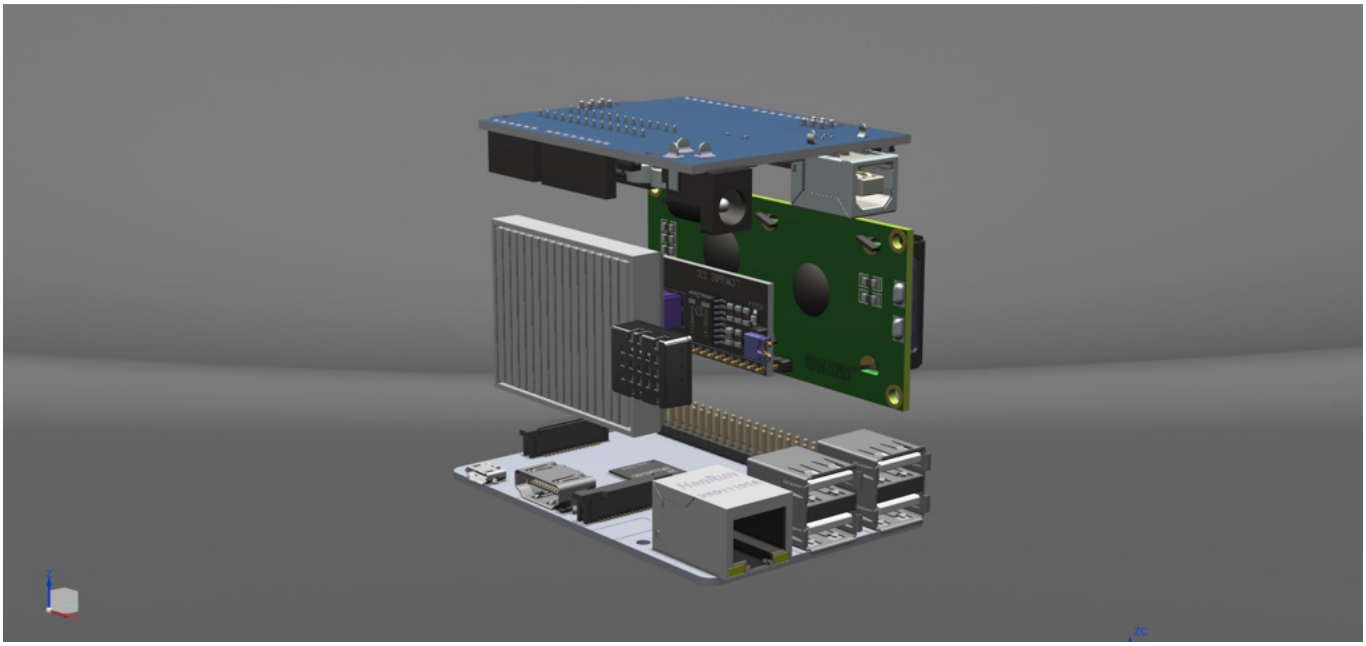

As shown in the 3D prototype above, A-Therm consists of a Arduino Uno module containing the DHT sensor, WiFi Module, Bluetooth Module, Micro USB cable, and LCD Monitor.

Final Product

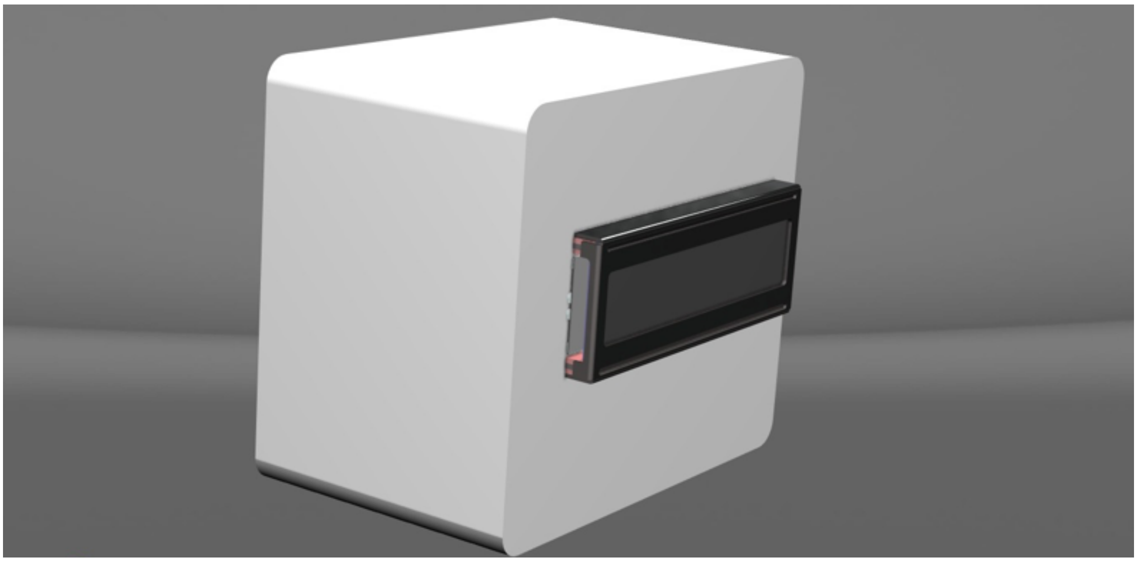

Isometric view of A-Therm and the LCD Monitor

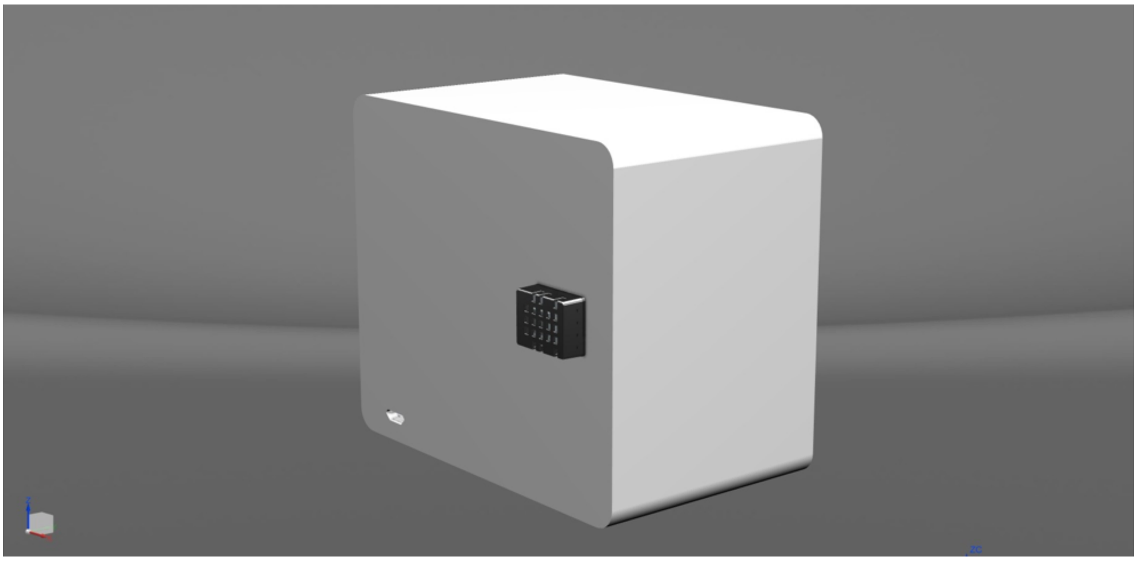

Isometric view of A-Therm and the DHT Sensor

Algorithm

A-Therm uses its own algorithm to calculate the optimal thermal comfort temperature using the outside temperature and humidity data.

The first step was to obtain a **realfeel** temperature using the outside temperature and humidity. A-Therm used **weather.gov**'s heat index calculator.

$$
\begin{aligned}
Index_{heat} = \; & -42.379 + (2.04901523 \times T) + (10.14333127 \times rh) \\
& - (0.22475541 \times T \times rh) - (6.83783 \times 10^{-3} \times T^2) \\
& - (5.481717 \times 10^{-2} \times rh^2) + (1.22874 \times 10^{-3} \times T^2 \times rh) \\
& + (8.5282 \times 10^{-4} \times T \times rh^2) - (1.99 \times 10^{-6} \times T^2 \times rh^2)
\end{aligned}
$$

In the above equation,
rh = relative humidity
T = actual air temperature

After obtaining the realfeel temperature from the raw data, A-Therm was to calibrate this value according to Human Thermal Comfort. For this, we incorporated *UC Berkeley's thermal comfort calibration calculator and grapher.* We input various values of realfeel temperature, as the grapher rendered a certain result each time.

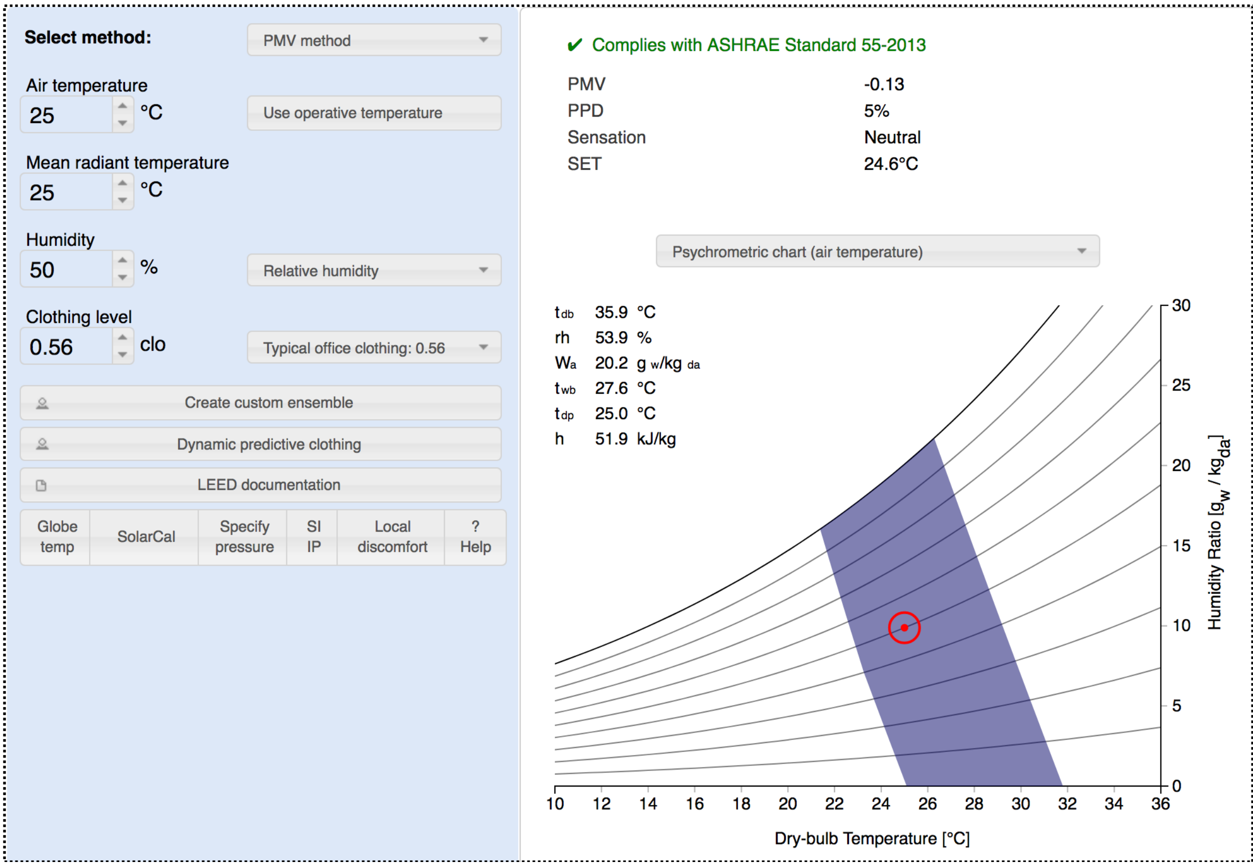

Plotting these results, we obtained a simple linear regression from the line of the best fit and used the equation in our code. The equation was as follows:

$$\text{TCT} = 14 + (\text{cel} * 0.3)$$

TCT = Thermal Comfort Temperature (Optimal Temperature)
cel = Actual Air Temperature (°C)

Although the above equation seemed to work just fine initially, it did result in inappropriate values of optimal temperature when the realfeel temperatures were around 18°C. To overcome this flaw, we modified the single equation into two **min. and max.** values.

$$TCT_{min} = 12.9 + (cel * 0.3)$$

$$TCT_{max} = 19.4 + (cel * 0.3)$$

TCT_{min} = Minimum Thermal Comfort Temperature
TCT_{max} = Maximum Thermal Comfort Temperature

From here, A-Therm calculated the average value of the two TCTs and therefore worked much more efficiently.

Ultimately, our calculations determined that A-Therm could save up to 30% electricity by **avoiding** the setting of the thermostat's temperature at mostly unreasonable levels. The Optimal Temperature ensures that the user gets as much as he/she needs to feel comfortable at any given weather conditions.

In comparison to a typical 3-star air-conditioner, A-Therm is almost 1.5 times more efficient in its energy saving ability.

Making the AI

One major problem we encountered was the difference of opinion among humans.

Indeed, the product rendered a certain value of Thermal Comfort Temperature (or Optimal Temperature) at which your body is *not supposed to* feel uncomfortable. But what if it doesn't suit your preference? What if you choose to have a hotter/colder temperature than what A-Therm suggested? Having to change the temperature on your thermostat every time the machine set it to something *you clearly do not like* would defeat the entire purpose of automation.

This led to our including artificial intelligence (AI) in the product.

The AI works in the following way

Let the user be suggested a value colder than the temperature at which he/she really feels comfortable (for example, 18 °C or 64.4 °F). Consequently, the user will change the value to a hotter level (for example, 21 °C or 69.8 °F).

A-Therm will **learn** this preference. For its next use, it will suggest a value comparatively hotter than its final result. How much hotter/cooler will be decided by the magnitude of difference between the user and machine's suggested temperatures. If it is not able to draw a clear conclusion about the preference, A-Therm will require changes in the suggested temperatures multiple times, until it is able to retain reliable information regarding the new setting.

Basic Mathematical Working

Let suggested Optimal Temperature = a ºC.
Let user's preferred temperature = b ºC.

Difference = (a − b) ºC.

Relative difference = $\frac{a - b}{a}$ = c.

New suggested temperature (for next time) = a(c+1) ºC.

*The above data is stored in the machine using **hash-tables**.*

Apart from the basic algorithm, we used the Google Cloud Machine Learning API to render the artificial intelligence more accurate.

A-Therm, therefore, learns something new about the user's preference every time it is used.

Key Outcomes

The research that was put into the creation of A-Therm unveiled a lot of important things. Upon basic observation, it can be discerned that most air conditioner consumers rely on a particular configuration to provide comfort in every weather.

Most often, it is the maximum cooling configuration which unnecessarily consumes more electricity and negatively affects the air conditioner's performance and life.

Based on mathematical hypothesis, ATherm was able to generate thermostat configurations that indicated more efficient cooling, lesser electricity consumption, and greater air conditioner longevity.

While researching for ATherm, we realized the importance ingrained in things like air temperature and humidity. We often fail to account for either while looking at it, assuming our comfort to be based only on one of the two parameters. ATherm fills this gap left by this shortcoming by accounting for it in its calculations.

Limitations

Despite its various advantages over a typical air-conditioner, A-Therm comes with its own drawbacks.

First, the device does not account for the number of people in the room. So, at a given thermal comfort temperature, the machine will not consider the carbon-dioxide levels in the room and will therefore not be accurate in party halls, theaters, and other 'crowded' areas.

Second, the AI currently used is quite simple and will need to be refined regularly in order to be accurate enough for the user experience. This would take months to implement correctly.

Third, the installation of the device requires the removal of the current thermostat an air-conditioner may use. Since A-Therm does not come in its own new air-conditioner, it may result in certain compatibility issues.

Last, but not the least, for some 5-star air-conditioners, A-Therm may not be significantly effective in reducing energy consumption. This makes it unappealing to invest in for AC's that are already very energy-efficient.

Future Work

A-Therm will gradually be refined toward something perfect after every few months. Below are some of our future plans regarding the device's functionality:

1. **Voice control**: using Google Cloud APIs, A-Therm can include voice commands in order to truly be a home automation product.

2. **Mobile app**: a mobile app (in the making) will serve not only as a remote, but also as an informative source for the amount of electricity saved every month. This way, the users can review their monthly usage and check whether or not the machine is truly benefitting them.

3. **Improvements in the AI**: the current AI is not perfect; with due passage of time, we plan to improve the machine-learning capabilities of the device.

4. **Combination Working**: for office cooling systems, A-Therm would be much more efficient if it were connected to a network a air-conditioners. This way, only one device would be required to control tens of air-conditioners, thereby cutting down on cost of investment.