Autonomous Sustainable Buildings: From Theory to Practice

Mario Bergés

2024-01-17

Table of contents

Preface

Buildings account for a large fraction of our total energy use. Any realistic decarbonization plan needs to consider them. Moreover, though technological upgrades (such as better insulation, newer and more efficient appliances, etc.) can improve their efficiency significantly, there is a relatively untapped potential for autonomous technologies to squeeze out much more out of the input energy sources feeding our building stock. This course intends to prepare students to become leaders in this nascent field that combines fundamentals of physics and building science with concepts from computer science, statistics and mechanical/electrical engineering, to name a few.

This website is built with Quarto. To learn more about Quarto books visit https://quarto.org/docs/books.

Part I

Preliminaries

1 How to use this book

> ⚠ Warning
>
> As stated elsewhere in the material, this book is still very much a work in progress. Nevertheless, here are (evolving) instructions for how to make the best use of it.

The material presented here is presented as both a website and a stand-alone book (in both PDF and EPUB formats). It can be daunting to know how it is all organized, so here are some tips for navigating the content. In particular, the type of information you may be seeking probably falls under one of these categories:

- Overview of what technical content is covered and in what order
- Lecture slides and notes
- Supplemental reading material
- Logistical and administrative information about the course
- Assignments and solutions
- Example problems

If the information you are looking for does not fall into any of those categories, please feel free to e-mail me. Otherwise, below is a more detailed description of each of those categories and where to find the corresponding content in this book.

1.1 Overview of technical content

Though the table of contents (shown on the left-hand side for the website, or in the initial pages for the PDF) is a good guide for the technical content contained herein, Chapter 3 contains a view of the schedule by which the content will be covered during class and may be a more useful guiding source. In particular, Chapter 3 has direct links to the lecture slides, the reference material and assignments for each lecture.

1.2 Lecture slides and notes

For each lecture, accompanying lecture notes are provided in the book. Additionally, during some lectures I also utilize slides and when I do, these are made available via this book as an external reference. To know which chapter of the book refers to a particular lecture, you will want to look at Chapter 3 (i.e., the schedule).

Lecture notes are currently very thin, as they are only meant to provide an orientation for how the lecture will proceed. However, they contain links to the slides (when appropriate), citations for the reading material from which the lecture was sourced, as well as a collection of questions that were left unanswered during the lecture.

1.3 Supplemental reading material

As stated in the previous section, supplemental reading material is typically provided as a reference for each of the lectures. To find this material you can either request a copy of the book or reference through the library or, more conveniently, navigate to the Files -> Readings section in Canvas to find a copy of the book chapters that were referenced.

1.4 Logistical and administrative information about the course

The full course syllabus is available on Canvas under Syllabus, or under File -> syllabus -> syllabus.pdf. An abridged version is also available in this book in Chapter 2.

1.5 Assignments and solutions

A list of assignments, along with their due date is available in Appendix B. It is also available in the course syllabus as found in Canvas.

1.6 Example problems

As we work on problems in class, I will make an effort to post those problems directly on the lecture notes for the corresponding lecture.

2 Syllabus

This syllabus will be continuously updated throughout the course.

2.1 Course Description

We spend a significant portion of our lives inside buildings: working, sleeping, and on leisure activities. Unsurprisingly, buildings are responsible for over 40% of our annual greenhouse gas emissions. The history of buildings goes hand in hand with the history of energy efficiency, as we have moved from non-renewable/inefficient fuel sources and technologies (e.g., firewood, cookstoves) to more renewable/efficient ones (e.g., solar energy, heat pumps). Increasing efficiency has also resulted in tighter integration between buildings, their systems, and the supporting services. All of these trends have led to an explosion in the number of instrumentation systems (for monitoring and/or controlling) installed in buildings, and an associated increase in the number and complexity of the decisions that are being (and can be) made in light of these new instruments. Autonomous technologies (which make decisions on our behalf in order to achieve pre-established goals) are well suited to address the challenges that these information-rich and highly-interconnected buildings pose.

With a focus on real-world deployments, case studies and group projects, this course will cover the theory and emerging practice of retrofitting existing buildings with hardware and software to significantly increase their autonomy and overall sustainability. The focus will be primarily on the operational stage of the life-cycle of buildings, and particularly on HVAC, electrical and water systems within them. In particular, this course will expose the students to recent advances in the quest to endow buildings with the ability to operate autonomously (i.e., with minimal human assistance), in order to provide the services they were designed for, while maintaining quality of service and upholding human values such as privacy, equity and environmental sustainability. The course will be based on lectures, assignments and a final project where the students will have the opportunity to design and implement an autonomous technology and evaluate it through a hardware-in-the-loop experiment in the lab. This course builds upon machine learning, statistics, data acquisition and instrumentation, linear systems and control theory.

It is intended to be a upper-graduate level course (i.e., for Ph.D. students or senior M.S. students) interested in gaining practical exposure to the state-of-the-art in data science for building systems. As such, the format of the course is tailored to that experience and will include reading and critiquing recent publications in the field, learning to implement data analysis techniques described in them, and producing novel results using these newly acquired skills.

The course assumes students are familiar with concepts in instrumentation, linear algebra, probability, statistics and programming, though this is not a strict requirement if the student has previously discussed with the instructor and has received approval.

2.2 Grading:

Here is how performance in the course will be evaluated:

Task	Percentage
Assignments	40%
Interim Project Progress Update	20%
Final Project Report and Demo	40%

2.2.1 Assignments

A total of four assignments will be given out. The topics covered in each assignment will closely follow the ones listed in the schedule of classes.

All assignments are to be solved individually. Discussions and conversations with other students regarding the problem sets are encouraged. However, the final solutions along with the reasoning behind them need to come from you and be clearly explained in the submitted documents.

Each assignment will be worth 10% and is due at the beginning of class on the date that is indicated in each assignment. Assignments that are submitted before this deadline can receive 100% of the available credit. There is a 3 day grace period for late submissions, with the following available credit in each of those five days: 1 Day Late: 90%, 2 Days late: 70%, 3 Days late: 40%. After this time, assignments will not be graded. Of course, if you anticipate not being able to meet this schedule due to a major problem, please talk to the instructor as soon as possible.

2.2.2 Project Updates

Another small portion (20%) of the grade for the course will be based on a progress update that will take place during the last period of the course. This update will consist of: (1) a written 2-page report, authored by all team members; and (2) a 10 minute individual meeting with the instructor to discuss the project, its goals and the plan forward. The written progress report will confer 10% of the final grade, while the individual meeting discussion will be used to provide the remaining 5%.

The final project report and demonstration are worth almost half of the total grade and, in some ways, are the most important assignment. The written report is worth 30% and the demonstration 10% of the total grade. For the written report, we are asking that you submit it using the ACM BuildSys Template (either LaTeX or MS Word). You should use your own words when writing it and avoid plagiarism of any kind. In it you will describe, in simple terms, the motivation and specific objectives of your project, the design choices made for the hardware and or software prototype that you put together, the experiments you performed to validate whether or not your solution satisfies the objective, and a discussion about the results and limitations. You should make all code and datasets available as part of your submission.

The rubric that will be used for grading the written report is as follows:

- Formatting and organization (15%)
- Grammar, clarity and accuracy of ideas (15%)
- Display of mastery of concepts covered in class (35%)
- Creativity expressed in the solution (10%)
- Depth of discussion related to the pros/cons of the implemented solution (15%)
- Overall assessment of the project's idea nd execution (10%)

2.3 Course Policies

Though there are definitely reasons to like anarchism, I still prefer the democratic system, so here are the rules of the game as they are now (and subject to change if enough of you request that I do).

2.3.1 Collaboration

Collaboration is expected within the limits of discussing concepts and problems. However, each student must produce his/her own solution to the problems. Copying from another student's assignment is clearly plagiarism. Using information directly from websites, books, papers and other literary sources without appropriate attribution is also plagiarism. Assignments submitted for this class will be reviewed by the instructor and TA and may be scanned through web-based academic integrity software. Occurrences of cheating or plagiarism will be handled according to the university policy on Academic Integrity, https://www.cmu.edu/policies/documents/Academic%20Integrity.htm. Students are expected to have read this policy and conform to the highest standards of academic integrity. For incidents of academic misconduct, the University Academic Disciplinary Actions Policy, found at https:

//www.cmu.edu/student-affairs/theword/acad_standards/creative/disciplinary.html, will be followed.

2.3.2 Class Participation

Students are expected to be in class on time and participate in class discussions. If you cannot make class, please inform your instructors and group members ahead of time. In class, students are expected to be courteous and respectful of the views and needs of other students and instructors.

2.3.3 Student with Disabilities

If you have a disability and have an accommodations letter from the Disability Resources office, I encourage you to discuss your accommodations and needs with me as early in the semester as possible. I will work with you to ensure that accommodations are provided as appropriate. If you suspect that you may have a disability and would benefit from accommodations but are not yet registered with the Office of Disability Resources, I encourage you to contact them at access@andrew.cmu.edu.

2.3.4 Posting of course materials

All the material used in the course (syllabus, readings, problem sets, reports) is intended for use in the class only. No unauthorized posting, publication or redistribution is expected. Uploading course materials to Course Hero or other web sites is not an authorized use of the course material.

3 Schedule

The schedule is subject to change. Here's the most updated version:

Date	Topic	Slides	References	Assignments
Part I	**FUNDAMENTAL CONCEPTS AND TOOLS**			
Tuesday 01/16/24	**1:** Introduction	Slides		
Thursday 01/18/24	**2:** Setting up your Computer Environment	Slides	Chapter 4	
Tuesday 01/24/24	**3:** Review of Data Acquisition and Instrumentation	Slides	Chapter 5	HW1 Out
Thursday 01/25/24	**4:** Setting up your Raspberry Pi SBC or Microcontroller		Chapter 6	
Tuesday 01/30/24	**5:** Energy Use In Buildings	on Canvas	Chapter 7; Pérez-Lombard, Ortiz, and Pout (2008); Chapter 1 from Murphy Jr (2021)	
Thursday 02/01/24	**6:** Building Thermodynamics: Part I and Part II	on Canvas		
Tuesday 02/6/24	**8:** Occupant Thermal Comfort	on Canvas		
Thursday 02/08/24	**9:** Basics of Energy Simulation: Part I	on Canvas		
Tuesday 02/13/24	**9:** Basics of Energy Simulation: Part II	on Canvas		
Thursday 02/15/24	**11:** Simulating and Estimating Occupancy	on Canvas		
Tuesday 02/20/24	**12:** Designing Controllers wioth Simulation Engines	on Canvas		
Thursday 02/22/24	**13:** Building Energy Simulation Wrappers: BOPTEST	on Canvas		

4 Setting up your computer environment

To better understand the contents of this course, we will encourage a hands-on approach to all of the learning activities. For this reason, preparing your computer for tinkering and playing around with these concepts is an important pre-requisite that we should get out of the way before diving into the fundamental concepts of building physics and other related topics.

In no particular order, here are the tools that we will be making use of throughout the course and that you should configure/install on your computer as soon as possible so that you can familiarize yourself with them:

- For programming:
  - Python 3, specifically the Anaconda Distribution
  - Jupyter Notebooks (already in Anaconda)
- For version control:
  - Git
  - A git client, such as the command line interface or a fancy graphical user interface.
  - GitLab, a self-hosted solution for interacting with Git repositories and adding some flare to them), in particular we will be interacting with this repository.
  - GitHub, another hosted solution for interacting with Git repositories. This is where you will host your final project.
- For Simulation:
  - EnergyPlus
  - BOPTEST
- For easy access and deployment of software solutions:
  - Docker
- For documentation and knowledge contributions
  - Quarto

During the lecture today, I will explore the use of **git**, **Quarto**, **Anaconda**, **Jupyter Notebooks** and **Python** all together by manipulating one particular notebook file found in this link:

- Lecture #2: Setting Up Your Environment [Jupyter Notebook]

The Jupyter Notebook in question (linked above) was generated directly from a source Markdown file using Quarto, which can ve found here: [Quarto Source]

Similarly, the same source file can generate other version of the same notebook such as: [HTML] and [PDF]

4.1 Following along

Here's a run down of the things we plan to do so that you can try them again at your own pace:

1) Install Anaconda following the instructions on the site.
2) Follow the Getting Started guide for Anaconda.
3) Use the command line through your comptuer's terminal emulator application (or the *Anaconda Prompt*) to test a few of the `conda` commands, including:
   - `conda --version` – To see the current version of *conda* that is installed.
   - `conda update conda` – To update *conda* to the latest version.
   - Work with environments by creating, activating and deactivating them:
     - `conda create --name [enviornment name] [packages to include]` to create an environment (substitute `[description]` with the appropriate value you want for it)
     - `conda activate [environment name]` to activate the environment
     - `conda deactivate` to deactivate and go back to the base environment
4) Work with Jupyter Notebooks and learn what they are. Specifically, we will follow the instructions on this Jupyter Notebook after downloading it and running `jupyter notebook` from the command line on the same folder where we downloaded it.
5) Make yourself familiar with the particular command-line shell used by your operating system. In particular, a Unix-like shell will be used extensively throuhgout the course and in the final project. As such, you may want to read some resources such as this or skim the contents of a full course such as this.
6) Work with *git* repositories. In particular, cloning, pulling, committing and pushing into a respository.
   - Create a public repository on GitHub
   - Clone the repository locally by issuing, on a terminal, the command `git clone [github repository's address]` (if you have SSH access configured, you can use the SSH address, otherwise use the HTTPS one).

- Add some files to the repository by, for example, copying the Jupyter Notebook we used earlier into the repository, and then staging it to be committed by issuing `git add Lecture_2_Setting_Up_Your_Environment.ipynb`
- Create a commit with a simple message by issuging `git commit -m 'Added a jupyter notebook'`
- To push the changes to the remote repository, you'll need to authenticate with GitHub. Nowadays, unless you use special tools such as the GitHub CLI, you will need to create a personal access token. You can do this by following these instructions.
- Using this access token, you can now proceed to push your changes by issuing `git push` or to be more precise `git push origin main` (which says, push to the remote repository called *origin* – now pointing to the remote server on GitHub – and, in there, to the branch called *main*).

7) Learn more about git by following the links on the Jupyter Notebook or this useful GitHub resource for it.

We will not discuss the details of BOPTEST, Docker or Quarto… but will do that later in future lectures.

4.2 What have you learned?

After the lecture, here are some questions that you may want to try to answer yourself to see if you really understood things.

Technical Questions

a) Assuming that I know the URL for the git repository that hosts the course's website, how do I go about cloning it?

b) What are branches in a git repository? Why are they useful?
c) After cloning a repository, is it advisable to make changes to it on the **main** branch, or on a separate branch created for that purpose?
d) What are conda environments and why are they useful?
e) How can I check the status of a local git repository from within a Jupyter Notebook?

5 Review of Data Acquisition and Instrumentation

Since for a building to be autonomous, it needs to be able to **sense**, plan, and act; then we need to review the principles of sensing in more detail. That is what we will do today.

5.1 Defining Autonomous Sustainable Buildings

Understanding the three terms in the title of this course (i.e., *Autonomous*, *Sustainable* and *Buildings*) is important before moving any further. In particular, the first two terms need to be defined ahead of time since buildings are a familiar object and we will spend a great deal discussing their inner workings in upcoming chapters. So what do we mean by a building that is *autonomous* and *sustainable*?

We can argue at length about the specific definition for these two terms so I want to start by pointing out that the specific definition is not that important (at least not right now). Instead, I find it more important for us to agree on the fact that sustainability is a goal, and that autonomous technologies are those that can achieve goals set by humans with minimal assistance from them. From this perspective, what we mean by the words in the title of the course is that its scope is defined by the technologies needed to allow buildings to achieve sustainable goals with minimal human intervention. Sustainability, of course, is a wide concept but Wikipedia's definition seems appropriate for us: "Sustainability is a societal goal that relates to the ability of people to safely co-exist on Earth over a long time". It is also worth pointing out that autonomy, in our view, is distinctly different from automation. While the latter refers to the ability of a system/process to execute some specific sequence of steps on its own, the latter refers to the ability of the system/process to not only enact those steps but also determine what they should be in the context of a given goal. This is, at least, our definition of these terms.

Autonomous systems follow a three-step cycle: sense, plan and act. Sensing refers to the ability of the system to perceive its environment and itself and develop some situational awareness. Planning (which in many contexts also includes learning), refers to the ability of the system to develop models of the world and itself based on the observations collected through sensing, as well as using these models to reason. Lastly, the actuation step allows the system to intervene and cause an effect in the environment, as decided upon by the planning step.

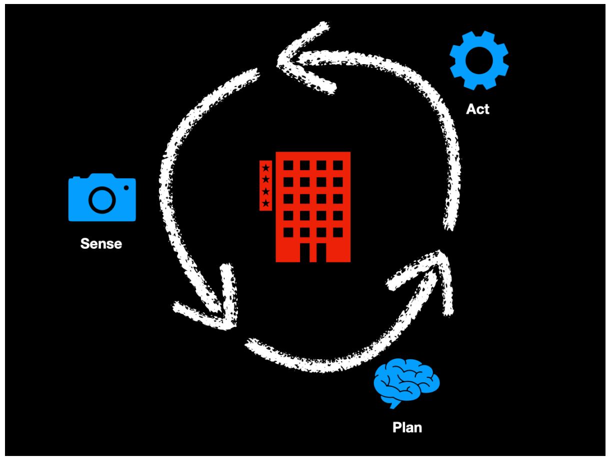

Figure 5.1: The autonomous loop present in all autonomous technologies: sense, plan and act.

5.2 Sensors

Given that our goal in the course is to create autonomous technologies to reach sustainability goals in buildings, we will need to focus our attention on this trifecta: sense, plan, act; and understand how it applies in the context of buildings. This chapter is intended to serve as a refresher for the main ideas behind modern sensing systems. For a more thorough treatment of this topic, we refer the reader to Fraden (2010), specifically Chapters 1 and 2. What follows is an abridged account of the main ideas needed to understand modern sensors and the process of turning analog signals into digital ones that can be processed by computers.

To start, let's take a look at some example sensing systems (which may do much more than just sensing) that have been developed in the last few years to tackle some building energy management problems:

Modern sensors and actuators are a relatively inexpensive commodity that has benefited tremendously from the technological advances of the past century to become cheap, miniaturized, reliable and massively produced. Most people living in the world today interact with sensors in one way or another and many carry multiple sensors in their pockets (e.g., on their mobile devices) all day. For our purposes, in this course, sensors are devices that can convert a physical stimulus into an electrical signal; actuators are devices that convert an electrical signal into a physical stimulus (the inverse of a sensor); and we can call either of these a transducer (i.e., a transducer is a term encompassing both sensors and actuators) and this term extends to include non-electrical devices too.

Given their wide availability and low-cost, there are a number of resources you can use to find sensors for your project in this course. Since we are going to be using Raspberry Pi computing platforms, here is a good starting point for finding sensing hardware:

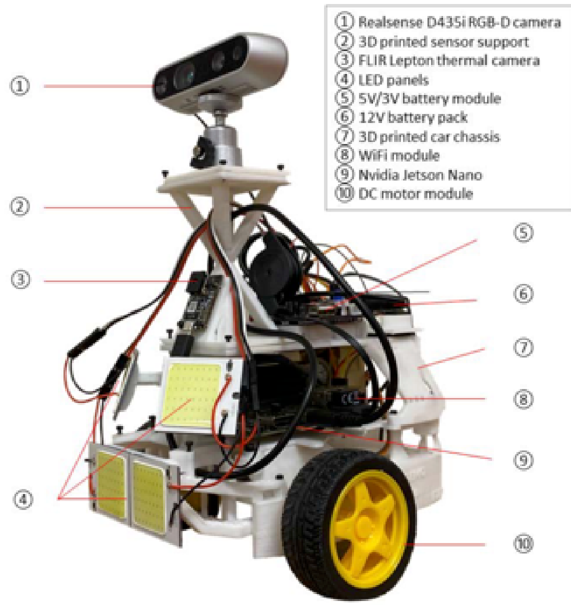

(a) TEA-bot Cai et al. (2022)

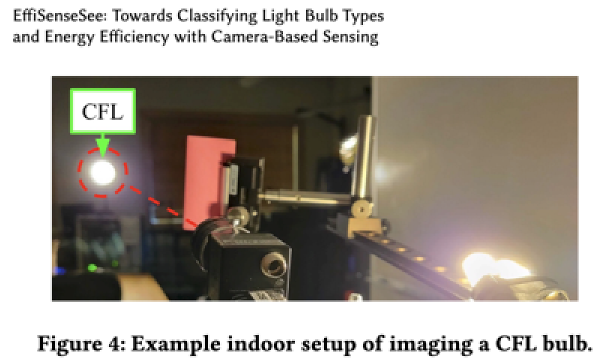

(b) EffiSenseSee Yen et al. (2022)

Figure 5.2: Some example sensing systems presented at the 2022 ACM BuildSys conference.

- Raspberry Pi Hardware by Adafruit
- Best Raspberry Pi HATs in 2023 by Tom's Hardware
- Raspberry Pi HATs, pHATS & GPIO by Pihut
- Rasbperry Pi Pico Hardware by Pimoroni

You should start thinking about what kinds of sensors you want to use for your project (if you know what its goal is), or get inspired by these sensors to define a goal for your project.

5.3 Digitizing Analog Signals

If sensors convert a physical stimulus into an electrical signal, how do we reason about that electrical signal to enable the planning and actuation phase of the autonomous solution we are designing? The answer, of course, is through the use of computers (broadly speaking, not just the electronic digital ones based on the von Neumann architecture we see everywhere today). That said, because our most successful approach at creating computers is based on digital electronics and the electrical signals that the sensors are producing will likely be analog, we need to involve (among other things) an analog-to-digital converter (ADC, or A/D converter) in the process. The topic of ADCs is wide and deep, and we are assuming in this course that you will have had some prior exposure to it. However, if you have not been exposed to it the minimum requirements for you to continue understanding the rest of the material are a basic understanding of the digitization process, in particular: sampling and quantization.

Since digital electronics process discrete information (bits), we need to discretize our incoming analog signal in both both time and magnitude. The process of discretizing time is called *sampling*, while the process of turning a continuous amplitude into a finite discrete set is called *quantization*.

5.3.1 Sampling

Let's start with the discretization of time. The interactive plot below shows a simple sine wave $x(t) = 1 + sin(2\pi \frac{f}{24} t)$ where f is the frequency of the sine wave and t is time measured in hours (the period of this signal is 24 hours, as you may have noticed). Let's imagine that this sinusoid is the physical phenomenon we are interested in measuring. This could be any property of the physical system we are twinning which, through a specific sensor, we are able to convert into an electrical signal. A simple example, in keeping with is that $x(t)$ represents the response of a photoresistive sensor sitting on the window as the sun goes up and down each day. If we take *snapshots*, or measurements of this signal at equally-spaced time intervals, the amount of time in between these measurements has a very important role to play in our ability to accurately represent the physical phenomenon we are measuring.

You can prove this to yourself by interacting with the simple simulation below.

```
//| echo: false

// Initial values
viewof frequency = Inputs.range([0.25, 10], { step: 0.25, value: 1, label: "Frequency (cycles
viewof samplingRate = Inputs.range([0.125,10], { step: 0.125, value: 1, label: "Sampling Rate
viewof showSineWave = Inputs.toggle({ label: "Show Continuous Sine Wave", value: true }); //

t = Array.from({ length: 2400 }, (_, i) => i * 0.01); // Time array for continuous plot from

// Continuous sine wave
continuous = t.map(t => ({ t, x: 1 + Math.sin(2 * Math.PI * (frequency / 24) * t) }));

// Sampled sine wave
sampled = Array.from({ length: Math.floor(24 * samplingRate) + 1 }, (_, i) => {
  const time = i / samplingRate;
  return { t: time, x: 1 + Math.sin(2 * Math.PI * (frequency / 24) * time) };
});

// Plotting
Plot.plot({
```

```
  x: { domain: [0, 24], label: "Time (hours)" },
  y: { label: "x(t)" },
  marks: [
    showSineWave ? Plot.line(continuous, { x: "t", y: "x", stroke: "blue" }) : null, // Condi
    Plot.dot(sampled, { x: "t", y: "x", stroke: "red", fill: "red" }) // Sampled points
  ]
})
```

Simply playing around with this can be very useful towards building intuition. But it helps to follow some directions, so you may want to wait until the activity at the end of this module for some specific guidance and suggestions for what to try.

5.3.2 Quantization

As we said initially, sampling is just one of the discretizations we need to do. But in addition to discretizing time, we also need to discretize the continuous amplitude of the analog signal through a process called quantization. Just as with sampling, quantization can introduce additional errors to our measurements. So it is worth paying close attention to how this process works in order to understand the possible errors. For instance, imagine that we continue trying to measure the intensity of light outside as driven by the day/night cycle on Earth. In the previous example we imagined that no matter when we measured the light intensity, our measurement apparatus would be able to provide us with the exact number representing the magnitude of light intensity at that moment. But what if we can only express that magnitude crudely, say by considering if it is “dark”, “dimly lit”, “bright” and “very bright” but nothing else? Well, it is obvious that this would be far from ideal to understand precisely how light intensity varies throughout the day.

Again, to build intuition regarding this concept, it helps to play around with the effects. The simulation below allows you to do just that:

```
//| echo: false

// Initial values for frequency, quantization levels, and sampling rate
viewof waveFrequency = Inputs.range([0.25, 10], { value: 1, step: 0.25, label: "Frequency (va
viewof quantLevels = Inputs.range([1, 16], { value: 4, step: 1, label: "Quantization Levels"
viewof newSamplingRate = Inputs.range([0.125, 10], { value: 1, step: 0.125, label: "Sampling
viewof newShowSineWave = Inputs.toggle({ label: "Show Continuous Sine Wave", value: true });

thirdTimeArray = Array.from({ length: 2400 }, (_, i) => i * 0.01); // Time array for continuo

// Continuous sine wave with a period of 24 hours
continuousWave = thirdTimeArray.map(t => ({ t, x: 1 + Math.sin(2 * Math.PI * (waveFrequency /
```

```
// Function to quantize the amplitude
function quantizeAmplitude(value, levels) {
  return Math.round((value * (levels - 1))) / (levels - 1);
}

// Sampled time points based on the sampling rate
sampledTimePoints = Array.from({ length: Math.floor(24 * newSamplingRate) + 1 }, (_, i) => i

// Quantized sine wave based on the sampling rate and number of quantization levels
quantizedWave = sampledTimePoints.map(t => ({
  t,
  x: quantizeAmplitude(1 + Math.sin(2 * Math.PI * (waveFrequency / 24) * t), quantLevels)
}));

// Generate grid lines for the background (for quantization levels)
quantGridLines = Array.from({ length: quantLevels * 2 }, (_, i) => i / (quantLevels - 1));

// Plotting
Plot.plot({
  x: { domain: [0, 24], label: "Time (hours)" },
  y: { domain: [0, 2], label: "x(t)", grid: true },
  marks: [
    newShowSineWave ? Plot.line(continuousWave, { x: "t", y: "x", stroke: "blue" }) : null, /
    Plot.dot(quantizedWave, { x: "t", y: "x", stroke: "red", fill: "red" }), // Sampled and q
    Plot.ruleY(quantGridLines, { stroke: "gray", strokeWidth: 0.5 }) // Grid lines for quanti
  ]
})
```

Activity: Exploring the Effects of Sampling and Quantization on Signal Accuracy

Use the interactive plot to experiment with signal sampling and quantization. Your goal is to understand how poorly chosen parameters can affect the measurement of a signal.
Instructions:

1) Start by hiding the continuous signal (toggle "Show Continuous Sine Wave" off).

2) Set the sliders to the following values for each case and observe the plotted points:

   - Case 1: Quantization levels = 8, Frequency = 6.75, Sampling rate = 1.
   - Case 2: Quantization levels = 2, Frequency = 11, Sampling rate = 1.
   - Case 3: Quantization levels = 16, Frequency = 10, Sampling rate = 0.375.

3) After exploring each case, answer the following:

- How do the sampled points differ from the original signal?
- How does reducing the quantization level or sampling rate affect the usefulness of the measurements?
- What are the key differences in signal behavior across the three cases?

```
//| echo: false

// Initial values for frequency, quantization levels, and sampling rate
viewof combinedFrequency = Inputs.range([1, 24], { value: 2, step: 0.25, label: "Frequency
viewof combinedQuantLevels = Inputs.range([1, 16], { value: 8, step: 1, label: "Quantizatio
viewof combinedSamplingRate = Inputs.range([0.125, 10], { value: 1, step: 0.125, label: "Sa
viewof combinedShowSineWave = Inputs.toggle({ label: "Show Continuous Sine Wave", value: fa

combinedTimeArray = Array.from({ length: 2400 }, (_, i) => i * 0.01); // Time array for con

// Continuous sine wave with a period of 24 hours
combinedContinuousWave = combinedTimeArray.map(t => ({ t, x: 1 + Math.sin(2 * Math.PI * (co

// Function to quantize the amplitude
function combinedQuantizeAmplitude(value, levels) {
  return Math.round((value * (levels - 1))) / (levels - 1);
}

// Sampled time points based on the sampling rate
combinedSampledTimePoints = Array.from({ length: Math.floor(24 * combinedSamplingRate) + 1

// Quantized sine wave based on the sampling rate and number of quantization levels
combinedQuantizedWave = combinedSampledTimePoints.map(t => ({
  t,
  x: combinedQuantizeAmplitude(1 + Math.sin(2 * Math.PI * (combinedFrequency / 24) * t), co
}));

// Generate grid lines for the background (for quantization levels)
combinedQuantGridLines = Array.from({ length: combinedQuantLevels * 2 }, (_, i) => i / (com

// Plotting
Plot.plot({
  x: { domain: [0, 24], label: "Time (hours)" },
  y: { domain: [0, 2], label: "x(t)", grid: true },
  marks: [
    combinedShowSineWave ? Plot.line(combinedContinuousWave, { x: "t", y: "x", stroke: "blu
    Plot.dot(combinedQuantizedWave, { x: "t", y: "x", stroke: "red", fill: "red" }), // Sam
    Plot.ruleY(combinedQuantGridLines, { stroke: "gray", strokeWidth: 0.5 }) // Grid lines
  ].filter(Boolean) // Filter out null marks
})
```

5.4 What have you learned?

After the lecture, here are some questions that you may want to try to answer yourself to see if you really understood things.

Technical Questions

a) According to the Nyquist-Shannon sampling theorem, if we want to sample a 100Hz signal, what sampling frequency should we use?

b) What is the resolution of a 16-bit ADC with a 10V full-rage voltage input?
c) If I sample a pure 100Hz signal at 60Hz, at what frequency does the resulting discrete signal repeat itself?
d) What are HATs in the context of Raspberry Pis and other embedded systems?
e) In the context of cars, which of the following is an example of an autonomous technology: cruise control, lane-keep assist, automatic parking.

6 Setting up your Raspberry Pi Devices

As we reviewed in the previous chapter, there are a myriad of sensors available out there that can help us enact the sensing infrastructure needed to create more autonomous buildings. These sensors, however, need to interface with some computing device before they can be useful to the second part of the sense, **plan** and act trifecta for any autonomous system. In this chapter, we will review and configure two platforms that can be used for this purpose, both designed and developed by the Raspberry Pi Foundation, namely:

- The Rasbperry Pi 4 Model B, from here on abbreviated as RPi 4.
- The Raspberry Pi Pico W, from here on abbreviated as RPi Pico.

The first of these, is a full single-board-computer (SBC) that has a 40-pin General-Purpose Input/Output (GPIO) port, with which one can communicate with digital transducers or with an external Analog-to-Digital (ADC) converter. The Pico W, on the other hand, is a board served by a custom microcontroller (the RP2040) that is significantly less powerful than the RPi 4 but also significantly less expensive and power-hungry. Depending on the application you have in mind, you may choose either (or both) of these platforms to interface with your sensors. Both of these devices have a wireless radio that allows them to connect to WiFi and be accessible remotely. Both can be programmed using Python (specifically through the help of MicroPython and CircuitPython).

The goal for today is to allow us to obtain measurements from a light intensity sensor (a photoresistor) using both the RPi 4 and the RPi Pico and, along the way, learn how to set up and configure these devices so that you are ready to use them when the time comes.

6.1 A Simple Photoresistor

The sensor that we will try to connect is this simple photoresistor:

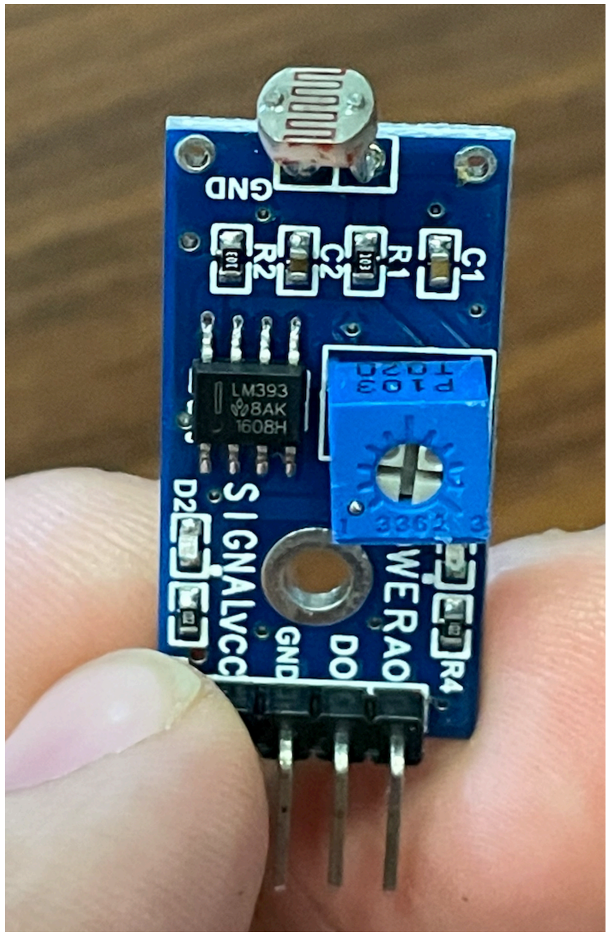

Figure 6.1: Picture of the photoresistor we will be using

This sensor reacts to light by decreasing its resistance through the principle of photoconductivity. The more light (luminosity) the lower the resisance, and with less light, resistance increases. We can measure the resistance by using Ohm's law ($V = IR$) and, assuming constant current, just measuring the voltage drop across this resistor.

As you can see in the figure, there is much more than just a photoresistor (the actual sensor is just the little thing sticking out on the top). The board also has an ADC, a trimmer potentiometer (trimpot), and other components. For this particular application, we are not interested in any of these and would just like to try to interface with the sensor as directly as possible (receiving the analog electrical voltage drop across the photoresistor). We can do this by focusing on two of the pins at the bottom: AO (Analog Output) and GND (Ground).

6.2 Raspberry Pi 4 Configuration

The RPi 4 does not come with an Analog-to-Digital converter that we can (easily) access. Thus, we'll need to interface it with an external ADC if we are to sample the raw analog output from the photoresistor. For this reason, before we can finish the task on the RPi 4 we will need to complete two steps: (a) set up the RPi so that it has an operating system (OS), an account we can access it with, and an Internet connection in case we want to access it remotely; and (b) use the GPIO port to interface with an ADC such as the MCP3008.

6.2.1 Setting up the RPi 4

To start, once you open up your RPi 4 kit, you will need to make sure that you have the following items in it (at least):

- A Rapsberry Pi 4 with heat sinks
- A mini-HDMI to HDMI cable
- A microSD card
- A microSD to USB adapter
- A USB-C cable

You will also need access to a USB power supply (your computer can perform that job), an HDMI monitor (at least to start), a keyboard, a mouse and another computer. Once the RPi 4 has been properly configured, you will only need the USB-C cable to power it.

Since these devices have been used before, a good starting point is to reflash the microSD card with a fresh install of the Raspberry Pi Operating System. To do that, follow the instructions here:

Install Raspberry Pi OS using Raspberry Pi Imager

In particular, you will download and load the Raspberry Pi imager, put the microSD card in the USB adapter and connect it to your computer, and follow the screen prompts on the imager to flash Raspberry Pi OS on the card (which will show on your computer in the same way that any externall mass storage device would). This process can take around 30 minutes or so depending on your computer and internet connection.

Once you do that, you can put the SD card back into the Rapsberry Pi's microSD reader, connect the keyboard, mouse and monitor to it, and power it via the USB-C port. You should be booting up to a new OS. If you installed an old version of the OS then the default username is `pi` with password `raspberry`. Otherwise, you will be prompted for a new password.

If it's all working correctly, you may want to take a moment to connect the device to the Internet. If you are at home, then follow the same procedure you would follow with any computer in your house. If you are on campus, you will need to connect it to the CMU-DEVICE network and, before that, you should register the device using the following website:

Device Registration. You will need to provide the MAC address for the Wireless card on your Raspberry Pi. This can be found in many ways, one of which is to issue `ifconfig` on a terminal in the RPi 4, or inspecting the contents of `/sys/class/net/eth0/address` (or `/sys/class/net/wlan0/address`).

6.2.2 Connecting the MCP3008

Let's follow the instructions here

Alternatively, we could follow instructions for a more modern package called `gpiozero` here.

6.2.3 Sampling from the photoresistor

The instructions above also include this part. Once the lecture is over, I will transcribe things here.

6.3 Raspberry Pi Pico Configuration

For the Pico, we will follow the instructions here.

6.4 Getting some help

Sometimes no matter how many references you have been provided, you can still get stuck on setting things up especially if it is your first time using these devices. Though I would strongly recommend against blidly trusting the guidance of a Large Language Model (LLM), if you are careful they can be of immense assistance to get you unstuck. For example, I asked ChatGPT o1 for assistance creating a guideline to this whole exercise and this is what it came up with: Raspberry Pi Light Sensing.

6.5 What have you learned?

After the lecture, here are some questions that you may want to try to answer yourself to see if you really understood things.

Technical Questions

a) What are the differences between the RPi Pico and RPi 4? Give an example of a project for which each one is better suited than the other.
b) We used a photoresistor for this example. Are the computing platforms (RPi 4 and RPi Pico) limited to photoresistors only? If not, what are the limitations we should consider when considering other sensors we could use with these platforms?
c) Does the RPi 4 have an analog-to-digital converter? If not, what is one way to sample analog signals using the RPi 4?
d) What is the clock speed for the RPi Pico and how do you find the answer using the `machine` MicroPython module? How is it related to the sampling rate we can achieve for sampling the analog signals with it?
e) What is the clock speed for the RPi 4, and how do you find out? Is it kept constant? How is it related to the sampling rate we can achieve for sampling analog signals?

Answers to Technical Questions

a) **Answered by:** *John*

The RPi 4 does not have a built-in ADC (analog to digital converter), so in order to sample analog signals on the RPi 4, we need an external ADC such as the MCP3008 we used in class. With this, we can connect the analog device to the external ADC, then communicate the converted digital signal to the RPi via another communication method (SPI, I2C, Serial, etc…).

7 Why buildings?

You have configured your computing environment, reviewed fundamental concepts in data acquisition and set up your Raspberry Pi devices. You are ready! But ready for what? In the previous chapter we discussed the three words on the title of this course (*Autonomous Sustainable Buildings*) and focused on defining the first two. But why are we targeting buildings in particular? Why not vehicles or financial trading agents? Well, of course, part of the answer is that you and I are both interested in buildings and in some sense that is sufficient. But are there good reasons beyond personal taste to focus our attention on buildings when developing sustainable autonomous technologies? Thankfully, the answer is yes. This last chapter of the preliminaries will provide the necessary context for all of us to answer that question more formally (or at least with hard evidence).

We will be relying on a few sources to understand the energy use of our building stock, namely:

- A paper by Pérez-Lombard, Ortiz, and Pout (2008) which discusses the building energy landscape as of 2008 using publicly available data at the moment.
- Chapter 1 from a new book by Murphy Jr (2021), which does some fun things to consider the growth of humanity's energy demand.
- Reports by the Energy Information Administration, especially the Commercial Building Energy Consumption Survey (CBECS) and the Residential Energy Consumption Survey (RECS).
- Lawrence Livermore National Laboratory's energy flow charts for energy use in the United States.
- Chapter 4 from Harvey (2010), which discusses more detailed statistics about energy use of bulidings worldwide.

Part II

Fundamental Concepts

8 Building Physics

We are now ready to address some of the fundamental technical concepts in building science, which are needed for us to reason about (*plan*) the data that our autonomous technologies will be gathering (i.e., what they *sense*). As we saw in the previous chapter, energy use in buildings is dominated by heating and cooling services. Thus, we will begin by reviewing basic concepts of thermodynamics leading up to the establishment of simple dynamic models of the temperature inside buildings and the energy required to maintain it.

8.1 Building Thermodynamics

In a very general sense, buildings loose (and gain) heat through interactions with the outdoor environment, and through internal processes that compensate for that. During a whole season (say, a full year), the energy lost through the building envelope is (again, at a high level) proportional to the leakiness of the building envelope, and to the temperature demand (i.e., the average difference between the indoor and outdoor temperatures). So, if we were to write an equation for the total (say, heating) losses Q that need to be compensated to keep the building at a given internal temperature (T_i), we can write the following high-level formula:

$$Q = \text{leakiness} \times \text{temperature-demand}$$

This leakiness is a function of the properties of the building's envelope. A bit more specifically, it is mostly affected by the way the envelope handles heat conduction and air infiltration and is usually expressed in units of energy per time per degree of temperature difference. The temperature demand is a function of the accumulated difference between the exterior temperature T_o and the average temperature we wish to keep the indoor space T_i over the season in question. This is usually described and tabulated as Heating Degree-Days (HDD) or Cooling Degree-Days (CDD), where Degree-Days refers to the accumulation (e.g., integral) of the difference in temperature ($T__o - T__i) over the season (expressed as number of days).

> Tip
>
> Read Appendix E of MacKay (2008) for more details on this topic. See if you can recalculate the numbers in Figure E.12 that refer to "my house, before" and "my house, after".

8.1.1 Heat transfer

Aside from air exchanges (e.g., infiltration), the leakiness defined above is caused by not just conduction but also convection and radiation. Together (conduction, convection and radiation) they form the basis for all modes of heat transfer. It is worth studying these in a little bit more detail.

8.1.1.1 Conduction

Joseph Fourier (1768 - 1830) came up with the law of conduction heat transfer, which relates the rate at which heat transfer occurs through a material, to the temperature difference it is subjected to, and the distance through which conduction is occurring. It is actually very similar to Ohm's law relating current (the rate at which electric charge flows through a conductor, to the voltage difference and the resistance). In particular, Fourier's law shows that:

$$\dot{Q} = -kA\frac{dT}{dx}$$

Where k is the thermal conductivity of the material, A is the area through which heat is flowing, and $\frac{dT}{dx}$ is the temperature *gradient* at a specific point x in the material.

Through a plane wall:

$$\dot{Q} = kA\frac{T_1 - T_2}{\Delta x}$$

(Assumming many things, including $T_1 > T_2$, no heat sources within the material, constant thermal conductivity, etc).

Which can also be stated as:

$$\dot{q} = \frac{T_1 - T_2}{\Delta x/(kA)}$$

Which is similar to how one would define Ohm's Law (i.e., $I = V/R$), and shows that the term $\Delta x/(kA)$ can be thought of as the resistance R of the wall.

2) Define fundamental concepts needed before proceeding futher:

- Temperature
- Heat Capacity: $C = cm$, $c = \frac{Q}{m\Delta T}$
- Heat transfer
  - Conduction
  - Radiation

  - Convection

3) Work on an example calculation of heat losses through a whole building envelope (problem set and spreadsheet on Canvas).

9 Thermal Comfort

Though so far we have spent time studying heat flows and thermal balancing through the building envelope, buildings are not designed and built just to shield ourselves from the external environment and its temperature. In fact, it can be argued that buildings are erected to improve the quality of life (or at least the productivity, in the case of commercial buildings) of its occupants. From that perspective, maintaining a proper temperature is but one of the many aspects required to achieve the greater objective. As such, the study of the quality of the indoor environment, including the acoustic, visual and thermal environment and its objective and subjective impact on the occupant is of great importance. This field is generally referred to as Indoor Environment Quality (IEQ), and includes sub-fields such as Indoor Air Quality (IAQ) that deal with explicit aspects of the environment.

In ways that are similar to how buildings regulate their internal temperature, our bodies do so as well. In particular, our bodies need to balance the heat that they generate (a byproduct of using up the chemical energy we ingest in the form of food) by dissipating excess heat to the environment and thus avoid overheating. This balance is the basis for human thermal comfort.

There are multiple mechanisms at play to achive this balance and self-regulate, but they all boil down to modifying the heat transfer rate from the body to the environment to match the rate at which heat is being transferred from the environment to the body. Though we have multiple such mechanisms (e.g., constricting blood vessels near the skin, sweating, shivering, etc.) they have limited capacity, as you may have discovered by venturing into the cold or hot outdoors without proper clothing. So it is no surprise that we need technologies (clothing, housing, etc.) to allow extend our bodies' self-regulation capabilities and be able to continue living in naturally harsh environments.

The total heat production rate for our body is the sum of the rate at which we produce work, and the rate at which we produce heat. Of these two, the heat production rate is the one that matters for thermal comfort. That said, there are many other external factors that affect comfort too including temperatuere, humidity, (solar) radiation and wind velocity.

The total heat ($\dot{Q}$) and work ($\dot{W}$) production rates together need to be equivalent to the rate at which heat is disspated through our skin. This equivalence is summarized in the following equation:

$$\dot{Q} + \dot{W} = \dot{M} A_{sk}$$

Where $\dot{M}$ is the metabolic rate, expressed in units of *met* (1 $met = 58.2\ W/m^2$), and A_{sk} is the total surface area of the skin. Of those values, the most important one for thermal comfort is $\dot{Q}$. Furthermore, A_{sk} is constant, so most of the change in this heat production rate is caused by the activities that the person is engaged in since they change the metabolic rate from, say, 0.7 *mets* for sleeping, to up to 250 *mets* or more when dancing or engaging in excersise. This rate of heat generated $\dot{Q}$ is then dissipated into the environment through the surface of the skin and clothing. Though the same kind of heat flow mechanisms we learned about before (namely: conduction, convection and radiation) are at play in this dissipation process, we should also take into account other heat flows due to evaporation (through skin) and other sensible and latent heat flows due to respiration. In other words, the major heat ransfer modes that account for the overall heat transfer rate $\dot{Q}$ are: $\dot{Q}_{\text{skin conduction}} + \dot{Q}_{\text{skin radiation}} + \dot{Q}_{\text{skin evaporation}} + \dot{Q}_{\text{respiration - sensible}} + \dot{Q}_{\text{respiration - latent}}$. Here in this chapter we will almost exclusively discuss radiation and convective (sensinble) heat flows through the skin and leave the evalporation and (sensible + latent) heat flows of respiration out of the explanation. The reader is encouraged to review Chapter 3 of Reddy et al. (2016).

Part III

State-of-the-Art Applications

10 Occupancy Estimation

This section will be populated soon...

Part IV

Implementation

11 Building a Novel Smart Electrical Meter

This document describes the motivation, objectives and initial ideas for Project #2 of the Spring 2023 edition of the course. The projects are the cornerstone of the course, as they are what motivates all of the material that is discussed during the semester.

The project's goal is to develop a computer system (hardware and software) that can monitor the electrical power usage of a section of the electrical power distribution system inside a building (e.g., at the main feed, sub-panel or individual circuit) and to leverage this information to support specific decisions by the building occupants or managers. These decisions need to be identified ahead of time by the project team, and can range from simple ones such as identifying base-load power consumption values, to more complex ones such as suggesting actionable steps to improve power quality or provide estimates of future power consumption under different scenarios.

There are various examples of projects (and commercially avaiable products) that can be used as a starting point. If you are curious to know more about them, a good place to start would be this (growing) list of links:

- EmonPi
- Split Single-Phase Energy Meter
  - Using the Split Single-Phase Energy Meter with RPi

12 Building a Novel Smart Thermostat

This document describes the motivation, objectives and initial ideas for Project #1 of the Spring 2023 edition of the course. The projects are the cornerstone of the course, as they are what motivates all of the material that is discussed during the semester.

The project's goal is to develop a computer system (hardware and software) that can interact with heating, ventilation and/or air conditioning (HVAC) equipment in a residential or commercial building, in order to enact "smarter" control of these, where "smart" here means that the new controller provides some cost, energy and/or productivity gains over the older one. One of the lowest-level controllers in modern HVAC systems consists of a (digital) thermostat and a Proportional Integral Derivative (PID) controller. Thus, the goal of this project is to create a computer system that can be used to replace an existing digital thermostat and increase the efficiency (energy, cost or productivity) of the building.

There are various examples of projects (and commercially avaiable products) that do just this. If you are curious to know more about them, a good place to start would be this (growing) list of links:

- ThermOS, with code here.
- DIY Thermostat with a Raspberry Pi
- Raspberry Pi as a touchscreen thermostat
- Building a Thermostat with the Raspberry Pi
- Programmable Thermostat with the Raspberry Pi
- PiTherm
- Raspberry Pi as a thermostat
- ThermTerm

A Useful Links

Throughout the course, we will encounter useful links that may not have an appropriate place within the structure of the website/book. This list will be the default placeholder for those links.

A.1 Shared News / Links

- Does turning the air conditioning off when you're not home actually save energy? Three engineers run the numbers
- Opportunities for students to work for and volunteer at Department of Energy affiliated sites
- Innovation in Buildings Graduate Research Fellowship (IBUILD)
- Atom Energy's solid state circuit breaker for EV charging
- Greta Thunberg's new book
- EV batteries alone could satisfy short-term grid storage demand by 2030
- An interesting learning resource for understanding Heating Ventliation and Air Conditioning systems
- Resources listed by the ACM Special Interest Group on Energy Systems and Informatics
- CMU's Energy Week event has a Student Happy Hour & Networking Reception on 3/21. Sign up here.
- An interesting discussion about heat pumps following the announcement of a Y-Combinator funded start-up trying to sell them directly to customers.
- ACM SIGEnergy's Graduate Student Seminar
- The Continental Automated Buildings Association (CABA) Podcast
- YouTube Channel TechnologyConnections covers thermostats

A.2 Interesting papers for second third

Here are some interesting papers to consider for the second third of the course, where we will be learning about the state-of-the-art in the field with respect to smart meters and smart thermostats.

A.2.1 Papers combining electrical meters and thermostat issues:

- **Non-Intrusive Techniques for Establishing Occupancy Related Energy Savings in Commercial Buildings**
- **Enhancing household-level load forecasts using daily load profile clustering**
- Comparing Gray Box Methods to Derive Building Properties from Smart Thermostat Data
- SMITE: Using Smart Meters to Infer the Thermal Efficiency of Residential Homes
- QUILT: QUantify, Infer and Label the Thermal Efficiency of Heating and Cooling Residential Homes
- Contextually Supervised Source Separation with Application to Energy Disaggregation

A.2.2 Papers mostly related to smart thermostats

- **ThermoCoach: Reducing Home Energy Consumption with Personalized Thermostat Recommendations**
- Operational Characteristics of Residential Air Conditioners with Temporally Granular Remote Thermographic Imaging
- Real-Time Cooling Power Attribution for Co-located Data Center Rooms with Distinct Temperatures
- Heat Reuse Models for Liquid Cooled Data Centers integrated with District Heating
- TEA-bot: A Thermography Enabled Autonomous Robot for Detecting Thermal Leaks of HVAC Systems in Ceilings
- The Impact of Resolution of Occupancy Data on Personal Comfort Model-Based HVAC Control Performance
- Good set-points make good neighbors - User seating and temperature control in uberized workspaces
- Hot or Not: Leveraging Mobile Devices for Ubiquitous Temperature Sensing
- Exploring Fairness in Participatory Thermal Comfort Control in Smart Buildings
- The SPOT* Personal Thermal Comfort System
- A toolkit for low-cost thermal comfort sensing

A.2.3 Papers mostly related to smart meters

- **Analyzing Energy Usage on a City-scale using Utility Smart Meters**
- BOLT: Energy Disaggregation by Online Binary Matrix Factorization of Current Waveforms
- Rimor: Towards Identifying Anomalous Appliances in Buildings

B Assignments for the Course

Assignment Number	Due Date	Resources
Assignment #1	2/6	[HTML] [PDF] [Jupyter Notebook]
Assignment #2	2/27	
Assignment #3	3/25	
Assignment #4	4/7	

Note: Solutions for all assignments will be posted on Canvas under Files -> Assignments -> Assignment sub-folder.

References

Cai, Weijia, Le Zhang, Lei Huang, Xinran Yu, and Zhengbo Zou. 2022. "TEA-Bot: A Thermography Enabled Autonomous Robot for Detecting Thermal Leaks of HVAC Systems in Ceilings." In *Proceedings of the 9th ACM International Conference on Systems for Energy-Efficient Buildings, Cities, and Transportation*, 30–39.

Fraden, Jacob. 2010. *Handbook of Modern Sensors: Physics, Designs, and Applications*. 4th ed. New York: Springer.

Harvey, Danny. 2010. *Energy and the New Reality 1: Energy Efficiency and the Demand for Energy Services*. Routledge. https://doi.org/10.4324/9781849774918.

MacKay, David JC. 2008. *Sustainable Energy–Without the Hot Air*. UIT cambridge. https://www.repository.cam.ac.uk/handle/1810/217849.

Murphy Jr, Thomas W. 2021. *Energy and Human Ambitions on a Finite Planet*. https://escholarship.org/uc/item/9js5291m.

Pérez-Lombard, Luis, José Ortiz, and Christine Pout. 2008. "A Review on Buildings Energy Consumption Information." *Energy and Buildings* 40 (3): 394–98.

Reddy, T, Jan F Kreider, Peter S Curtiss, and Ari Rabl. 2016. *Heating and Cooling of Buildings: Principles and Practice of Energy Efficient Design*. CRC press.

Yen, Alex, Zeal Shah, Benjamin Ochoa, Pat Pannuto, and Jay Taneja. 2022. "EffiSenseSee: Towards Classifying Light Bulb Types and Energy Efficiency with Camera-Based Sensing." In *Proceedings of the 9th ACM International Conference on Systems for Energy-Efficient Buildings, Cities, and Transportation*, 69–78.