

# CS397/497, Spring 2024

## Final Design Project

This is an **individual** assignment. You will be creating a writeup which can be in any clean format you choose, but please clearly label sections of the writeup. When finished, submit a PDF through Gradescope.

### Introduction

The purpose of this assignment is to gain understanding of the power profile, communications capability, economics, and other key tradeoffs for existing and emerging IoT connectivity options in real-world application contexts. The emphasis is on design space exploration. For the Cellular homework, we gave you a specific application scenario (these many bits, this often, etc). In real-world engineering, however, your job is to explore the tradeoff space—e.g., perhaps your application could sacrifice latency, send messages twice as long half as often, and still work “well enough”?

Your task is to prepare a report on the design space options. Materially, this will look like a lot of different, simple graphs (e.g., x-axis bytes/day, y-axis \$/day). Several of these will probably be straight lines. The most interesting ones are the ones with “knees” or other inflection points (e.g., the first 1MB data/day is cheap, then gets pricey).

**The goal is to gain experience in synthesis of comprehensive system design. This is not expected to be a perfect analysis, rather an order-of-magnitude estimate that you would provide to upper management during the early design phase to guide major project direction.**

#### Facility Maintenance Technician III Salary in San Diego, California

[Change City](#)

How much does a Facility Maintenance Technician III make in San Diego, CA? The average Facility Maintenance Technician III salary in San Diego, CA is **\$69,653** as of October 27, 2022, but the range typically falls between **\$62,941** and **\$77,652**. Salary ranges can vary widely depending on many important factors, including [education](#), certifications, additional skills, [the number of years](#) you have spent in your profession.

What do I mean by this? We ask for an estimate of labor costs to deploy sensors. Make an educated guess and cite a source. E.g. if I were deploying sensors in a building in San Diego, I might Google “San Diego building maintenance salary” and cite/use the first [reasonable-looking result](#). Then  $\$70k / (50 \text{ weeks} * 40 \text{ hours}) = \$35/\text{hr}$ . Not quite as abstract as a [Fermi Estimate](#), but pushing towards that line of thinking. Find reasonable numbers, justify them, and move on. Give

a sense of the space of options for your chosen application. The emphasis here is on **quantitative** over qualitative evaluations, *every claim and consideration should have numbers that justify it.*

Throughout your report **cite your sources**. *Make it easy to see where your numbers came from.* Include both the URL and a screenshot of the specific part of the resource you used to get your numbers.

There is no explicit page count: it's more important that you actually cover the topics listed on the next several pages. A concise paper that meets the requirements gets a better grade than a long paper which ignores them. A really good paper has clear indications of which parts of the paper are for which requirements with bolded headings. Papers are usually in the 6–15 page range, with the larger page count often due to images, graphs, and tables (not just text).

## Part A: Defining Your Application

Your first task will be to define your application. Have some fun with this and use your imagination. Try to come up with a real-world use case for an “IoT” system. We will impose a few constraints on your application:

- 1. Your application must be “large-scale,” defined as 100+ individual sensor nodes deployed.**
  - These may be densely deployed, e.g. in a ‘smart building’ context, or spread out over a wide geographic area, e.g. in an ‘environmental monitoring’ context, or whatever is appropriate to the proposed application. The constraint is simply that you have a large number of devices.
- 2. Your nodes must have at least two operating modes with different communication requirements.**
  - The two modes could be for different purposes, or different states of one purpose. Lots of reasons why they might exist that you can play with.
  - More modes are possible, as makes sense for the application.
- 3. Your edge nodes cannot be plugged in.**
  - If you propose supporting infrastructure, e.g. border routers, hotspots, etc., those may be powered.
- 4. Your application must be deployed in a specific location.**
  - It may be the same country that you chose for your cellular exploration, or a different one. There’s no limit on the size of the area, as long as it’s specific.
- 5. Your application must run for (at least) one full calendar year.**

To help you brainstorm, we list examples of real-world applications in [Appendix A](#).

### [45 pts] Part A Deliverables

There are not precise, specific values here. The goal is to establish a range of “what is reasonable” for your application. This will set the bounds for the design space exploration later.

[10 pts]: Give a high-level description of your application.

- What will your system measure, and why is that data important?
- Who are the stakeholders who would be interested in the data your system collects?
- What is the physical scale needed for your application—geographic area and number of sensors?
- How does the application meet all of the above requirements? Don’t forget that there must be two operating modes.

[10 pts]: What constraints are there on the data your system collects for the stakeholders, e.g.:

- Latency/timeliness? (Realtime? Daily updates? Hourly?)
- Network reliability/robustness? (How much data can be lost? From individual devices, the system as a whole?)
- Volume/Rate? (How much data and how quickly?)

[10 pts]: Without doing any numerical analysis (yet), rank the wireless technologies we have discussed in class from best to worst for your proposed application.

- Options should include:
  - BLE
  - 802.15.4/Thread
  - WiFi
  - Legacy Cellular (2g/3g)
  - IoT Cellular (LTE Cat-M / NB-IoT)
  - LoRaWAN
- If you believe that another option we talked about in class fits well, you may consider it.
- Qualitatively, you should consider things such as infrastructure needs / availability, network topology (and how that related to deployment), costs to incorporate the technology into devices, and costs to operate the technology. Give a ***brief*** justification for your rankings. Make sure your justification includes all options above.

***For the rest of this assignment, you will only consider the top two technologies chosen here.***

[10 pts]: Given the two wireless technology candidates, what constraints are there on the deployment of your system and how do they apply to the communication modes, e.g.:

- Infrastructure? (Is there existing infrastructure you will need to use/pay for? Or will you need to have a solution for how you will deploy infrastructure, e.g. putting up 15.4 border routers in a building?)
- Cost? (What is the order-of-magnitude target cost for the whole system; what does that come down to per-device? Consider both infrastructure costs and network usage costs, as applicable.)
- Deployability / Maintenance (How much time will it take to deploy devices; if you need to access them again [e.g. to change batteries...] how much time will that take; what is the (rough) cost of labor?)

[5 pts]: What is another constraint you think is important to consider for your application?

## Part B: Energy Modeling

We have mentioned several times that a primary reason IoT devices use different wireless technologies is due to resource constraints, of which energy is critical. We haven't done much to quantify that yet, however.

One task you have is to figure out how much energy your sensor will spend on communication over the lifetime of its operation. To make this a bit easier, we're just going to consider the energy spent on communication related activities, that is, for the energy analysis you can assume that the platform does no sensing and has no application to run other than just waking and sending data when ready. You should consider the major energy needs for each mode of operation of your sensor node. For example:

- Energy/latency to power on, off?
  - If you turn off the radio between events
- Energy consumed between events?
  - If you leave the radio on (in sleep mode?) between events
- Energy during uplink communication
  - For example, data-send events
- Energy during downlink communication
  - If any exist
  - For example, firmware-update events

### Some References

Where to find power numbers? Datasheets! Here are some examples that might help out. You're free to choose different radio hardware.

BLE/802.15.4/Thread:

- nRF52840 microcontroller, see section 6.20.15 (Electrical Specification for the Radio peripheral): [https://infocenter.nordicsemi.com/pdf/nRF52840\\_PS\\_v1.7.pdf](https://infocenter.nordicsemi.com/pdf/nRF52840_PS_v1.7.pdf)

WiFi:

- ESP32-S3 microcontroller, see section 4.6 (Current Consumption): [https://www.mouser.com/datasheet/2/891/esp32\\_s3\\_datasheet\\_en-2946743.pdf](https://www.mouser.com/datasheet/2/891/esp32_s3_datasheet_en-2946743.pdf)

Cellular:

- One chipset you might choose, check out section 4.2.3 (Current Consumption): [https://www.mouser.com/pdfDocs/SARA-R5\\_DataSheet\\_UBX-19016638.pdf](https://www.mouser.com/pdfDocs/SARA-R5_DataSheet_UBX-19016638.pdf)
- Empirical analysis of power modes on real hardware: <https://www.digikey.com/en/articles/how-to-enable-power-saving-modes-of-nb-iot-and-cat-m>
- Release 13 Spec [Includes CAT M1, CAT NB1]: [https://www.etsi.org/deliver/etsi\\_ts/124000\\_124099/124008/13.07.00\\_60/ts\\_124008v130700p.pdf](https://www.etsi.org/deliver/etsi_ts/124000_124099/124008/13.07.00_60/ts_124008v130700p.pdf)

- Paper providing a “power consumption analysis of NB-IoT and LTE-M”  
<https://ieeexplore.ieee.org/document/8644481>

LoRa:

- SX1262 transceiver, see section 3.5.1 (Power Consumption):  
[https://www.mouser.com/datasheet/2/761/DS\\_SX1261-2\\_V1.1-1307803.pdf](https://www.mouser.com/datasheet/2/761/DS_SX1261-2_V1.1-1307803.pdf)

Datasheets, like specifications, are huge and cover every little detail. You probably only need the information from one Table in one page of the thousands here. Use the Table of Contents and Search to guide you. For energy, sections titled “Electrical Characteristics/Specifications” or “Typical Operating Behavior” are usually where you will find the table you will want.

Quick EE Refresher: Most datasheets list current draw of operating modes. They also list the ‘nominal operating voltage’. The last piece of the energy puzzle is ‘how long am I in this operating mode’. E.g., if a radio draws 10 mA when sending, operates at 3.3 V, and is active for 1 second, then  $10 \text{ mA} \cdot 1 \text{ s} \cdot 3.3 \text{ V} \approx 33 \text{ mJ}$  energy / packet.

Note: You don’t need to do a comprehensive search to find the lowest power radio for your technology. Just pick one and use its values. Do a search for “Low Power LoRa Radio” or similar and that is good enough.

## [20 pts] Part B Deliverables

[10 pts]: Present the model of the energy use of communication for your application. This should include multiple radio states and account for application communication patterns including data amount. Ensure that this is a **quantitative** model with real numbers for cited sources.

[5 pts]: Based on the model, compare the expected energy performance for the two wireless technologies selected in Part A that you are studying. Your analysis should answer, “in what cases is technology A better and in what cases is technology B better”?

[5 pts]: As a final step, what is one example of a consumer-grade battery solution (i.e. 1 AA, 3 AAA’s, a D cell, 2 CR2032’s, etc) that would be a reasonable choice (cost, form factor...) to power each case? You may use a simple battery capacity model—every battery chemistry has a nominal voltage, and batteries are rated in mAh of capacity. The old cell phone battery (lithium polymer, 3.7 V nominal) on my desk is rated for 1800 mAh, so it has  $1800 \text{ mAh} \cdot 3.7 \text{ V} \approx 24 \text{ kJ}$  of energy. Again, provide numbers and justification for *why* this battery solution would suffice.

## Part C: Making Your Case

In Part A, you developed a list of application constraints. Now, your job is to apply those constraints to the two wireless technologies you are considering.

### [35 pts] Part C Deliverables

[25 pts] Present a comparison of the expected performance for the two wireless technologies you are studying based on the constraints determined in Part A. You do not necessarily need to present graphs or tables for every single constraint, only the ones that are most relevant for comparison given your application and your technology. *Convince a fellow engineer or scientist that your analysis covers the important tradeoffs.*

- You must compare on at least three different metrics.
- Some comparisons may result in proof that a certain metric is not relevant for this application, but at least some metrics will be relevant.

[10 pts] Give a final assessment on how you think your application should be built given the tradeoffs you have analyzed. Which technology should you use, how should it be deployed, what kind of data should it send, how often should it send it, and what will it cost to keep your deployment alive for one full year?

## Appendix A: Real-World IoT Deployments

These are just some examples to inspire you, lifted from [this survey paper](#):

The “Internet of Things” describes a wide and diverse range of applications. To understand and quantify their networking requirements, here is a survey notable application papers from the sensor networking literature and their networking requirements in two deployment scenarios.

The first, single location case, assumes the application is deployed to the fullest extent in a single location. We report the throughput and range required to support these deployments by multiplying the number of nodes in the deployment by the amount of data per measurement by the sampling interval.

A single instance of an application is often not consistent with the ubiquity targeted by the IoT. Therefore, we also consider the pervasive case for each application, which assumes that the application is scaled to be fully deployed in its target environment. For example, while a single location case may describe an application that monitors a single building, the pervasive case would include monitoring for all buildings of that type throughout a city. The applications vary tremendously in deployment area, so we employ the bit flux metric to compare them in terms of bits per hour per square meter.

The networking requirements for the eleven applications surveyed are shown in the table, and are described below, along with the assumptions for their pervasive deployments.

<b>Application</b>	<b>Single Location Throughput (bps)</b>	<b>Single Location Radius (m)</b>	<b>“Bit Flux” (bits-per-hour/m<sup>2</sup>)</b>
<b>Zebranet</b>	53	75	0
<b>Trash Can Monitoring</b>	0.38	370	0.003
<b>Hospital Clinic</b>	11	20	0.02
<b>Volcano Monitoring</b>	520	1,500	0.2
<b>CitySee</b>	20,400	5,700	1
<b>Electricity Metering</b>	51,389	6,180	1.5
<b>Habitat Monitoring</b>	10	10	9
<b>H1N1</b>	18,000	60	43
<b>IMT-2020</b>	33,556	564	128
<b>Macroscope</b>	12	4	221
<b>GreenOrbs</b>	5,600	80	1,000



**Zebranet** is one of the earliest sensor network research deployments. It places GPS tracking collars on zebras that asynchronously send location data over a wide-area network. The incredibly low density of wild Grevy's Zebras results in near zero bit flux over a wide area, with peak throughput coming from monitoring all zebras in a large herd.

Pei Zhang, Christopher M. Sadler, Stephen A. Lyon, and Margaret Martonosi. 2004. Hardware Design Experiences in ZebraNet. In Proceedings of the 2nd International Conference on Embedded Networked Sensor Systems. 227–238.

**Trash Can Monitoring** reports when trash cans are full in a deployment of 197 monitored trash cans throughout New York City's Times Square. Each trash can reports approximately twice a day, and we assume the same frequency and density for a pervasive deployment.

Bigbelly, Inc. 2019. New York City's Times Square Efficiently Manages 26,056 Gallons of Waste and Recycling Each Day with Bigbelly. <http://info.bigbelly.com/case-study/times-square-new-york-city>

**Hospital clinic** measures patient vital signs in a 32 bed hospital clinic in St. Louis, USA. At scale all patients in the 2915 hospital beds in St. Louis would be monitored.

Octav Chipara, Chenyang Lu, Thomas C. Bailey, and Gruiua-Catalin Roman. 2010. Reliable Clinical Monitoring Using Wireless Sensor Networks: Experiences in a Step-down Hospital Unit. In Proceedings of the 8th ACM Conference on Embedded Networked Sensor Systems. ACM, 155–168.

**Volcano monitoring** senses seismic tremors across 16 devices on Reventador volcano in Ecuador, streaming data when an event is detected. The pervasive case covers a volcanic area at the same sensor density.

Geoff Werner-Allen, Konrad Lorincz, Jeff Johnson, Jonathan Lees, and Matt Welsh. 2006. Fidelity and Yield in a Volcano Monitoring Sensor Network. In Proceedings of the 7th Symposium on Operating Systems Design and Implementation. USENIX Association, 381–396

**CitySee** measures air quality from 1196 devices deployed in Wuxi, China, and we assume the same sensor density for a pervasive deployment.

Xufei Mao, Xin Miao, Yuan He, Xiang-Yang Li, and Yunhao Liu. 2012. CitySee: Urban CO<sub>2</sub> Monitoring with Sensors. In IEEE International Conference on Computer Communications (INFOCOM'12). 1611–1619

**Electricity metering** in San Francisco, USA. Approximately 370,000 smart meters throughout the city report 250 byte readings once every four hours.

Dominic Fracassa. 2019. CleanPowerSF tripling households served with municipal electricity. <https://www.sfchronicle.com/bayarea/article/CleanPowerSF-tripling-households-served-with-13618155.php>. Pacific Gas and Electric Company. 2016. EPIC 1.14 - Next Generation SmartMeter Telecom Network Functionalities.

**Habitat monitoring** measures microclimate and occupancy of bird burrows with 32 sensors on Great Duck Island off the coast of Maine, USA. A pervasive deployment would monitor the estimated 5000 Storm Petrel nests on Great Duck Island with 7500 sensors.

Alan Mainwaring, David Culler, Joseph Polastre, Robert Szewczyk, and John Anderson. 2002. Wireless Sensor Networks for Habitat Monitoring. In Proceedings of the 1st ACM international workshop on Wireless sensor networks and applications. Acm, 88–97. College of the Atlantic. 2018. Seabirds at Great Duck Island. <https://www.coa.edu/islands/great-duck-island/seabirds-at-gdi/>.

**H1N1** measures a single-day human contact graph of 850 people for modeling flu epidemiology in a school in San Francisco, USA. A full deployment would measure interactions for the 80,000 students in San Francisco.

Maria A. Kazandjieva, Jung Woo Lee, Marcel Salathé, Marcus W. Feldman, James H. Jones, and Philip Levis. 2010. Experiences in Measuring a Human Contact Network for Epidemiology Research. In Proceedings of the 6th Workshop on Hot Topics in Embedded Networked Sensors. ACM.

**IMT-2020** defines performance characteristics of 5G technologies. For machine-type communications it defines a connection density of one million devices per km<sup>2</sup> each transmitting a 32 byte packet every two hours.

ITU-R. 2017. Guidelines for Evaluation of Radio Interface Technologies for IMT-2020.

ITU-R. 2017. Minimum Requirements Related to Technical Performance for IMT-2020 Radio Interface(s).

**MacroScope** monitors the microclimate of a redwood tree with 33 sensors on a tree in Sonoma, USA. A full deployment would place sensors on all trees in an old-growth forest, at a density of about 20 trees/acre.

Gilman Tolle, Joseph Polastre, Robert Szewczyk, David Culler, Neil Turner, Kevin Tu, Stephen Burgess, Todd Dawson, Phil Buonadonna, David Gay, et al. 2005. A MacroScope in the Redwoods. In Proceedings of the 3rd International Conference on Embedded Networked Sensor Systems. ACM, 51–63.

National Park Service. 2015. Del Norte Coast Redwoods.

<https://www.parks.ca.gov/pages/414/files/DelNorteSPFinalWebLayout2015.pdf>.

**GreenOrbs** measures ecological data from 330 devices in a forest near Tianmu Mountain in China. We assume pervasive deployment at the same sensor density.

Lufeng Mo, Yuan He, Yunhao Liu, Jizhong Zhao, Shao-Jie Tang, Xiang-Yang Li, and Guojun Dai. 2009. Canopy Closure Estimates with GreenOrbs: Sustainable Sensing in the Forest. In Proceedings of the 7<sup>th</sup> ACM Conference on Embedded Networked Sensor Systems (SenSys '09). 99–112.