

# Smart Library

Authors: Krish Vaswani, Pablo Wilson, Arjun Raguram,  
Electrical and Computer Engineering, Carnegie Mellon  
University

**Abstract**—In the time of coronavirus, social distancing has become an important aspect of human life. Public spaces that once had space for plenty of people must now limit their intake because of social distancing. In order to organize the availability of different spaces, we introduce Smart Library. In this project, we set up a system of cameras in a public space which constantly monitor the availability of socially distant seating and transmit this information to a public website. This is achieved with four cameras sending information to a central node, which preprocesses the data and sends it to a server in the cloud to perform image recognition and parse the information. The information is then transmitted to the public through a website.

**Index Terms**—Camera, Sensor, Machine Learning

## I. INTRODUCTION

Libraries are a huge part of the lifestyle of a university student. It is a haven where students can go to do homework or focused studying for an upcoming exam. On the Carnegie Mellon University campus, one of these libraries is Sorrell's, which can get very crowded. Finding seats there can be difficult, even when you take into account social distancing protocols during a pandemic. An existing solution is a project called Carrel Corral, which detects the carrels available in Hunt Library using a system of sensors and LEDs that send information to an app. One problem with design is that it physically interferes with the space of the library with wired components, which could impede the work of a student. Also, this design only works in one place (Hunt Library).

Smart Library is a system that aims to help college students find proper seating in a library or public space through wireless methods. It consists of a wireless system of cameras connected to microprocessors that will take bird's-eye view images of a space and send them to a cloud server to run image detection of seats available for use. It also aims to analyze daily/weekly behaviors of a library of when seats tend to be available. The users of the Smart Library system will access the information obtained via a website. Our main goal is to use 4 images from 4 cameras to send it and run image detection on a cloud, as well as send the information to the website in less than or equal to 1 minute. This will translate to a camera frame rate of 1 frame per second. We will also aim to achieve a battery life of 72 hours and an object detection accuracy of 76.5%.

## II. DESIGN REQUIREMENTS

Since the project was inspired by the lack of availability of

space at Sorrell's Library at Carnegie Mellon University, our design requirements were derived from this one use case. By observation, we know the seats are usually occupied in the span of a few minutes after being left vacant. Therefore, it is necessary to update information at least once per minute. Our design choices were mainly based on this time constraint. Additionally, the funds acquired for this project was \$600, and some design decisions were made to remain under budget.

Besides the time and budget constraints, there were two other minor constraints. If the batteries need to be replaced too often, it would become inconvenient for the people managing the library, and therefore Smart Library would become a burden to maintain. In order to ensure that the project remains relevant, we require the system to be able to function independently for 72 hours without changing the battery. Additionally, we have a requirement for the accuracy of image recognition for the project. The image recognition is based on the YOLO model, which has a top-1 accuracy of 76%. In our implementation, we modify this model to detect specific objects. Since this is more specialized, we also require the implementation to be at least as accurate as the original model that is more general.

Component	Time Allowance for Task Completion
Camera Node	20s
Central Node	30s
Cloud	10s

Table 1

In order to ensure the design meets the time specifications, the requirements are further divided into three parts. Namely, camera nodes, the central node and the cloud. The breakdown of the timings for each section is given in Table 1. These timings are based on the specific function and hardware chosen for each portion, which are described in the Architecture and System Description sections.

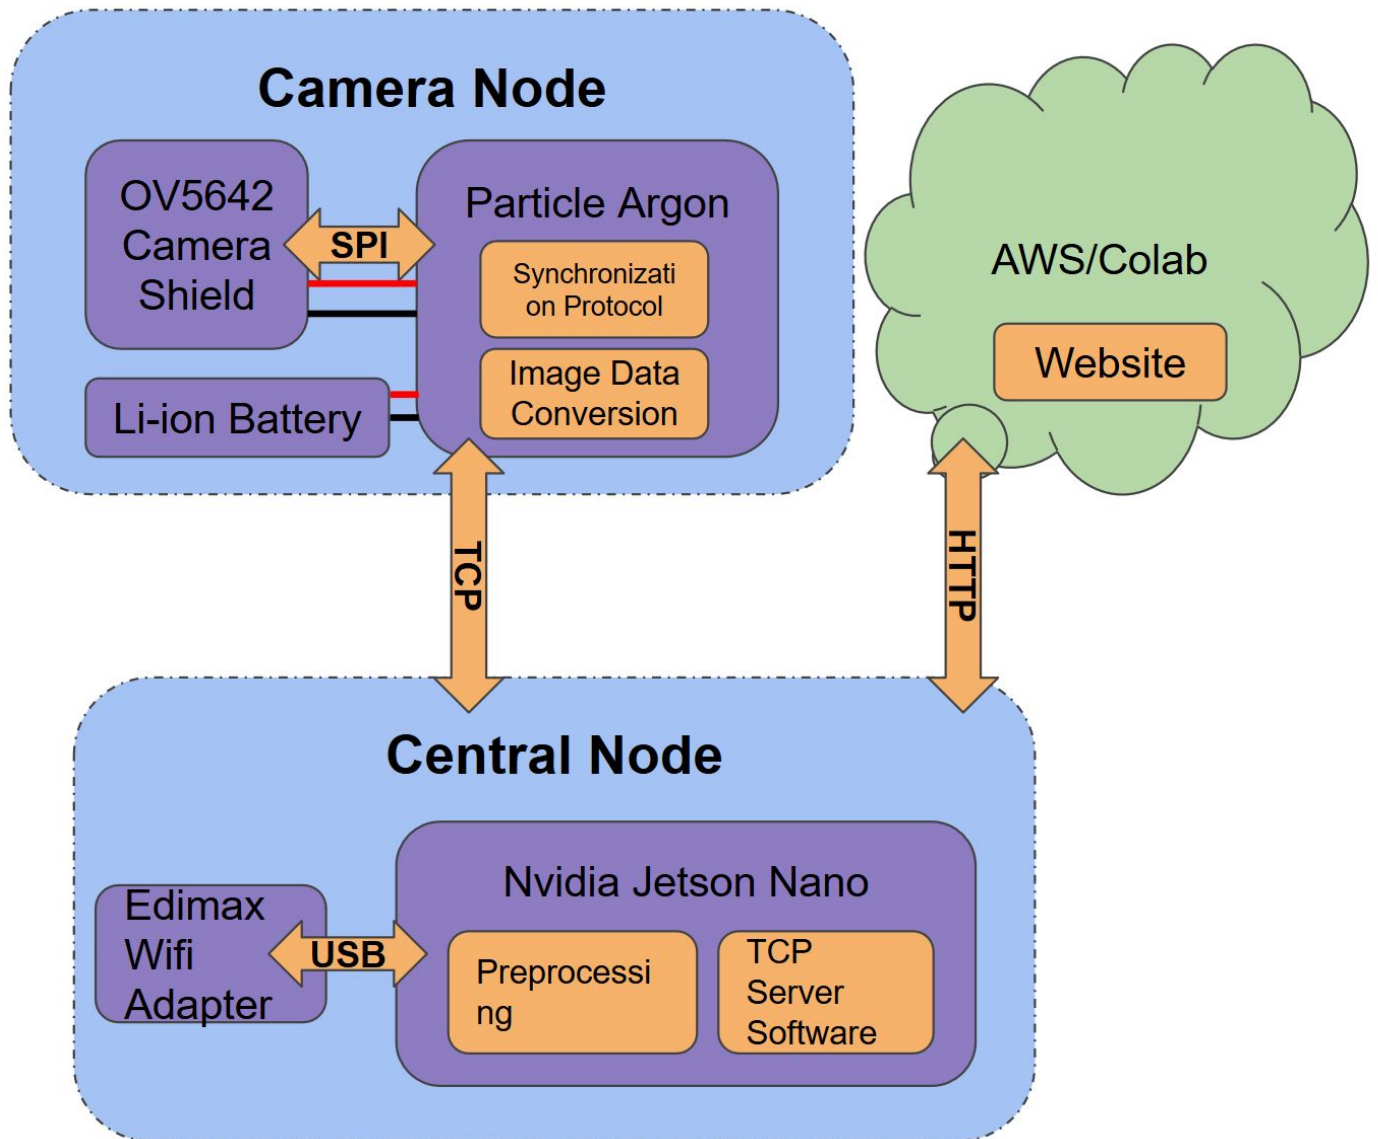


Fig. 1. Block Diagram of Node Network with 1 Camera Node

### III. DESIGN TRADE STUDIES

#### A. Camera Node

One design trade off that was made for the Camera Node was for the camera module. Ideally, we would like the highest resolution image possible, but we had to balance this with both cost and time to upload to the Central Node. We chose the OV5642 Camera Shield because it fit within our budget for each Camera Node and because the 5MP resolution it provides allows us to capture sufficiently high resolution images for the Object Detection Model and be able to upload the image over Wi-Fi at fast enough speeds to meet the upload time requirement. Another design trade off we made was choosing an MCU base for each Node. We decided on the Particle Argon because it comes with an integrated Wi-Fi chip and has sufficient processing power for the work it needs to do.

#### B. Central Node

One design trade off that was made for the central node was using an NVIDIA Jetson Nano instead of using another device like Raspberry Pi or INTEL NUC or NVIDIA TX2. One of the reasons is because of the use of image processing on the central node. As we work with the object detection algorithm in the cloud, we will have to determine what sort of image data pre-processing is needed to present the data to the model in an efficient and correct way. Image processing is generally faster on an NVIDIA Jetson Nano than on a Raspberry Pi, so any uncertainty with processing power or time will be accounted for here. Another reason for choosing the Jetson Nano over other NVIDIA devices like the TX2 or Jetson Xavier was due to a combination of portability and cost. The Jetson Nano is a relatively small computing device that costs around \$99, which is much cheaper than the other NVIDIA alternatives. We only had a \$600 budget in total, so getting something more expensive would not be ideal especially if we needed replacements.

The main reason we chose the Edimax 2-in-1 Wifi and Bluetooth Adapter is because there has been previous projects and records of it being used with a Jetson Nano and working properly. When it comes to using external devices, the computer needs to have the correct drivers available if not installed already with the operating system. Since the NVIDIA Jetson Nano does not have the same kind of processor as a personal laptop or desktop, the operating system and existing drivers may not be compatible, so it is important to use an external Wifi adapter where it is more or less confirmed to work by another individual. The Wifi adapter also has enough bandwidth and network speed for our purposes too, which was a necessity.

#### C. Cloud computing

One of the design trade offs that we needed to make for cloud computing was the size of the instance. AWS instances

usually come in xlarge, 2xlarge, 4xlarge and 8xlarge sizes, with each one being roughly two times faster and more expensive than the last. In our preliminary research, we found that the machine learning algorithm chosen was designed to run quickly on a small graphics processing unit (GPU). So, the smallest instance was sufficient. Additionally, in our application, we would need to constantly run the model over an extended period of time. The cheapest option allows us to run the model for the longest time given a fixed budget. Therefore, the cheapest option, which is xlarge in this case, is the most optimal.

### IV. SYSTEM DESCRIPTION

As mentioned previously, the system is divided into three parts, camera nodes, a central node and the cloud. The camera nodes will take pictures periodically and send them to the central node, which will aggregate the pictures, perform preprocessing and send the final product to the cloud. Then, object detection will be run on the cloud to parse information, and display it on the website.

#### A. Camera Nodes

Each Camera Node will consist of three hardware components: Particle Argon, OV5642 Camera Shield, and a 3.7V 1800mAh rechargeable LiPo battery. We decided to use the Particle Argon primarily due to its integrated Wi-Fi module, capable of reaching speeds up to 100Mbps, more than enough to meet the design requirements for upload speed. There will also be two software components: a synchronization protocol and image data conversion. Each Camera Node will be mounted high on a ceiling or wall to provide a bird's eye view.

The main purpose of each Camera Node is to capture images of the area from a bird's eye view and upload them to the Central Node. The Particle Argon will be programmed with a synchronization protocol capable of synchronously taking images with the other Camera Nodes, so that images are taken and fed to the Central Node with minimal discrepancy. The Particle Argon will also have data conversion programming to convert the image data from the camera module to a form that the Central Node wants to receive.

The other purpose of each Camera Node is to survive at least 72 hours on its battery life. The synchronization protocol will be designed to keep the Camera Node active only when needed to prolong the battery life.

#### B. Central Node

The central node will consist of two hardware components:

NVIDIA Jetson Nano and Edimax 2-in-1 Wifi and Bluetooth Adapter (EW-7611ULB). We decided to use the Jetson Nano instead of another device like a Raspberry Pi because running high-computational image processing we would use before sending the images to the cloud would have a better performance output on a Nano. The Edimax 2-in-1 is a known adapter to work well with the Jetson Nano drivers and has a network speed of 26 Mbps, which is enough for the project since we require 5-6 Mbps to adhere to the design requirements.

The main purpose of the central node is to receive packets wirelessly from camera nodes and send packets wirelessly to the cloud. The Nano will have server-side socket programming that will receive TCP packets synchronously and concurrently from the four camera nodes. The reason that TCP was chosen over UDP is because with TCP, packet loss is not an issue. Even though UDP offers a faster alternative to TCP, there still lies the risk of packet loss, and we require the central node to take a total of 30s of processing time, the speed of TCP will suffice. For concurrency, the use of pthreads will be used to parse incoming packets from four camera nodes, and since the processor on the Nano is quad-core, there will be one thread running on each core, thus a performance advantage. We will be designing our own network application protocol to enable proper synchronous communication between the central node and camera nodes. The central node will communicate with the cloud and send data packets via HTTP requests since the cloud hosts the website as well.

Another purpose of the central node is to conduct image pre-processing on packets before sending image data to the cloud. This includes any pre-processing that would make the data presentation to the YOLO model (discussed in section C) sufficient as well as image stitching. Image stitching is an algorithm that takes multiple images and connects them together to create one image that acts as a map of an entire area. For the purpose of Smart Library, it will be used to combine 4 images into one overall map of a public space. This will help make the AWS model run smoothly.

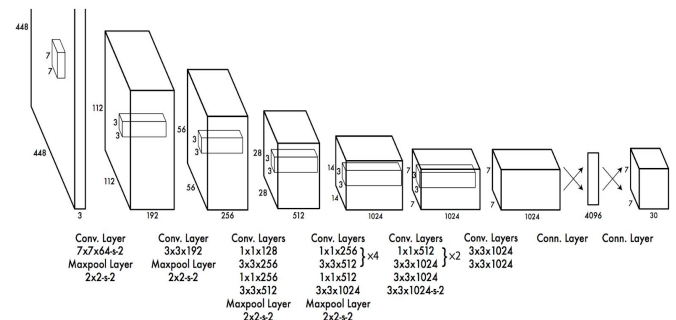
### C. Cloud

The cloud will run two pieces of software, an object detection model and a website. Cloud computing was chosen because it enables renting computing resources. This is especially helpful with machine learning models like the one we use in this project. Cloud computing gives us cheap access to a Graphics Processing Unit (GPU), which is required for machine learning.

The specific cloud computing resource we chose for this project is Amazon Web Services (AWS) and Google Colab. Google Colab offers free access to a GPU, making it good for development of the machine learning model. However, Colab is also inconsistent in giving its resources, so we will have to move the project to AWS for the final stages of development

and for deployment. AWS offers access to an NVIDIA T4 GPU with the g4dn.xlarge hardware instance. This instance costs \$0.526 per hour and allows us over 175 hours on the budget allocated for cloud computing, which is more than what is required for the project. Additionally, g4dn.xlarge has a fast network speed of 25 gigabits per second, which will reduce latency when communicating with the other parts of the project. For software, the Ubuntu Deep Learning AMI will be used, because it comes pre configured with common deep learning libraries like PyTorch and OpenCV.

The cloud will run a machine learning algorithm which will be adapted from the YOLO model [1]. YOLO is a convolutional neural network designed for real-time object detection. This particular model was chosen because of the time constraint. Other state of the art object detection models are built for non-real time applications, and therefore sacrifice latency for accuracy. Since time is a more important metric in this project, we decided to use the YOLO model.



The YOLO architecture is depicted in the above figure. It contains seven convolutional layers followed by two fully connected layers and post-processing. For this project, we will maintain the architecture and weights of the first few layers and change the last few layers. The original model was trained to identify 80 classes. For our purpose, however, we need to identify only two classes, vacant and occupied seats. The first few layers of the model will process the image and output an abstract representation from which it is easy to extract information about the contents of the image. The model then processes this representation to find out which of the 80 classes are found in the image, along with their locations. For this project, we will take the abstract representation and train a custom model to process this information to find out which of our two classes are found in the image, along with the locations. The new architecture will be determined after data has been collected, since the number of parameters and the quality of parameters is dependent on the quality and variety of data available.

Once the locations are determined by the object detection model, the data will be abstracted into a map to protect the privacy of people in the public space. This map will be publicly accessible through a simple website that only processes get requests, which will be implemented in Django.

## V. PROJECT MANAGEMENT

### A. *Schedule*

Our Schedule contains primary tasks groups: Camera Nodes, Central Node, Object Detection Model, and Website. By the nature of the project, the hardware must be complete before the Object Detection Model can be properly trained, thus the Camera Nodes and Central Node task groups with be tackled first, concurrently. After Camera Nodes has achieved its milestones, work on the Object Detection Model task group will begin in full force. The Website task group is designed so that tasks may be worked on between gaps in other task groups. All task groups are scheduled to be completed by the first week of December.

We encountered setbacks due to shipping delays of our components since the initial phases of our project relies heavily on having working nodes to provide our ML algorithm trainable data. Since our original schedule was aggressive, with a final milestone scheduled for the end of November, our timeline shift still has us finishing by the end of the semester.

### B. *Team Member Responsibilities*

Each team member will take primary responsibility for a task group. Pablo will take lead on the Camera Nodes, Arjun will take lead on the Central Node, and Krish will take lead on the Object Detection Model. Secondary responsibilities include integration between task groups and the Website task group. Pablo and Arjun will work together to develop the communication between the Camera Nodes and the Central Node. Arjun and Krish will work together to develop the image pre-processing and transferal to the Object Detection Model from the Central Node. Krish and Pablo will work together to find best image capturing practices to feed to the Object Detection Model. Everyone will work on the Website in the gaps between Task Groups. For Pablo and Arjun, this is anticipated to be after their task groups are complete. For Krish, this is anticipated to be before the bulk of his task group begins.

### C. *Budget*

Of the roughly \$600 we received, our budget breakdown is as follows: Camera Nodes, \$100 each, 4 total nodes; Central Node, \$100 each, 1 total node; AWS credits, \$100, 200 hours total. The specific costs breakdown are as follows: The Camera Nodes each require 3 primary components, the Particle Argon, \$25; Camera Module, \$50; and Battery, \$10. The Central Node requires 2 components, the Jetson Nano, \$100; and the WiFi adapter, \$15. AWS credits are paid for each hour used at a rate of \$.526 per hour.

### D. *Risk Management*

The way we mitigated design risk was to have a generic system architecture with clearly defined needed features. No

single aspect of our design is locked to a specific component; our components were chosen solely due to the fact they could achieve the features we wanted for the cheapest price and/or easiest integration. For example, our Central Node can be replaced by any hardware component(s) that can connect to the internet and do image processing and our Camera Nodes can be replaced by any hardware component(s) that can capture images and transmit them to the Central Node. A goal of our design was to be able to be as flexible as possible, so that it should be near arbitrary to scale up or swap components.

The way we planned for schedule risk was by setting ourselves with an aggressive timeline so that we would finish early if everything went to plan. We mitigated risk of unanticipated delays this way by allowing our schedule to shift by a couple weeks and still be on track to finish before the deadline. Additionally, we split our tasks into groups so that we can easily adapt and adjust when any one task gets behind schedule, as well as having primary task groups be staggered, so that there is always a team member available to assist if needed.

We managed resource risk by ordering duplicate parts and avoiding purchasing all parts at once. Even though our design for a fully functioning node network contains at least four Camera Nodes, our design will still function with just one, albeit with less area covered. Furthermore, full functionality can be tested with just two Camera Nodes, so we initially only purchased enough components for two Nodes in case we ran into a design problem and had to pivot to a different design for our Camera Nodes. For our Central Node, due to the high cost of the Jetson Nano, we were not able to purchase duplicate nodes. However, we only purchased one of the Wifi and Bluetooth Adapter, so that Arjun could confirm the part worked for our design before Pablo purchased one to be put into the Node Network on campus.

## VI. RELATED WORK

One related project was done in Fall 2019 of CMU Capstone 18500 by Team A4. The project was called Carrel Corral. The main goal of this project was to find available study carrels on the third floor of Hunt Library on CMU campus. Their approach was to have a system of ultrasonic and infrared sensors at each carrel connected to a microprocessor to determine whether there was a person there, as well as have a mobile app that can communicate this information to possible student users. They also considered the case where a student places their backpack at a carrel but does not actually sit there for a long period of time.

## VII. SUMMARY

A. *Future work*

Although this project is part of a semester-long course, we might continue working on it after the semester ends. Throughout the semester, we plan on demonstrating the project in Sorrell's Library at Carnegie Mellon University. However, this can be extended to any general space. One particular type of space is restaurants, who might want to monitor seating availability in order to automatically assign tables without the need for a host. This would require modifying the system to fit into multiple different environments. Specifically for object detection, data would have to be collected at a much larger scale from multiple locations in order to build a model that can generalize more easily.

Another modification that could be made is integration with existing hardware. It is common for public spaces to have security cameras. For installation in such a space, we could remove the requirement for our own cameras in lieu of the security cameras that have already been installed. This would require a change in hardware along with compatibility with the specific security system of the space. The advantage with this method is that our system would be less invasive since it reduces the requirement of installing additional hardware.

B. *Lessons Learned*

[TBD]

## REFERENCES

- [1] You Only Look Once: Unified, Real-Time Object Detection, <https://arxiv.org/abs/1506.02640v5>
- [2] National Semiconductor Inc., [www.national.com](http://www.national.com).
- [3] CMU 18500 Capstone Team A4: Carrel Corral, <http://course.ece.cmu.edu/~ece500/projects/f19-teama4/proposal/>

## IX: Schedule Diagram

