Carnegie Mellon

Electrical & Computer ENGINEERING

Search and Shine

David Wu, Nina Guo, Ronit Shah

Electrical and Computer Engineering Department, Carnegie Mellon University

http://www.ece.cmu.edu/~ece500/projects/S24-teamc7

Introduction

In the complex and hazardous environments of disaster stricken areas, traditional Search and Rescue (SAR) rescue methods expose workers to several dangers and inefficiencies. Recognizing these challenges, we have designed our **Search and Shine rover, an autonomous vehicle that uses computer vision to detect humans and report real-time information to a secure website,** enabling remote coordination with rescue workers.

Our team defined several use-case requirements for the Search and Shine rover. It must **accurately point a laser** at a human to demonstrate human tracking. It also must have **longevity in autonomous drive** whilst searching for humans. Real-time updates are also extremely important in SAR missions, so there must be **fast response times** to the web application. We have designed a system that meets these important requirements using **Computer Vision** on multiple processors, **Control logic** to move the rover and camera autonomously, and **low latency communication over WiFi**.

System Architecture

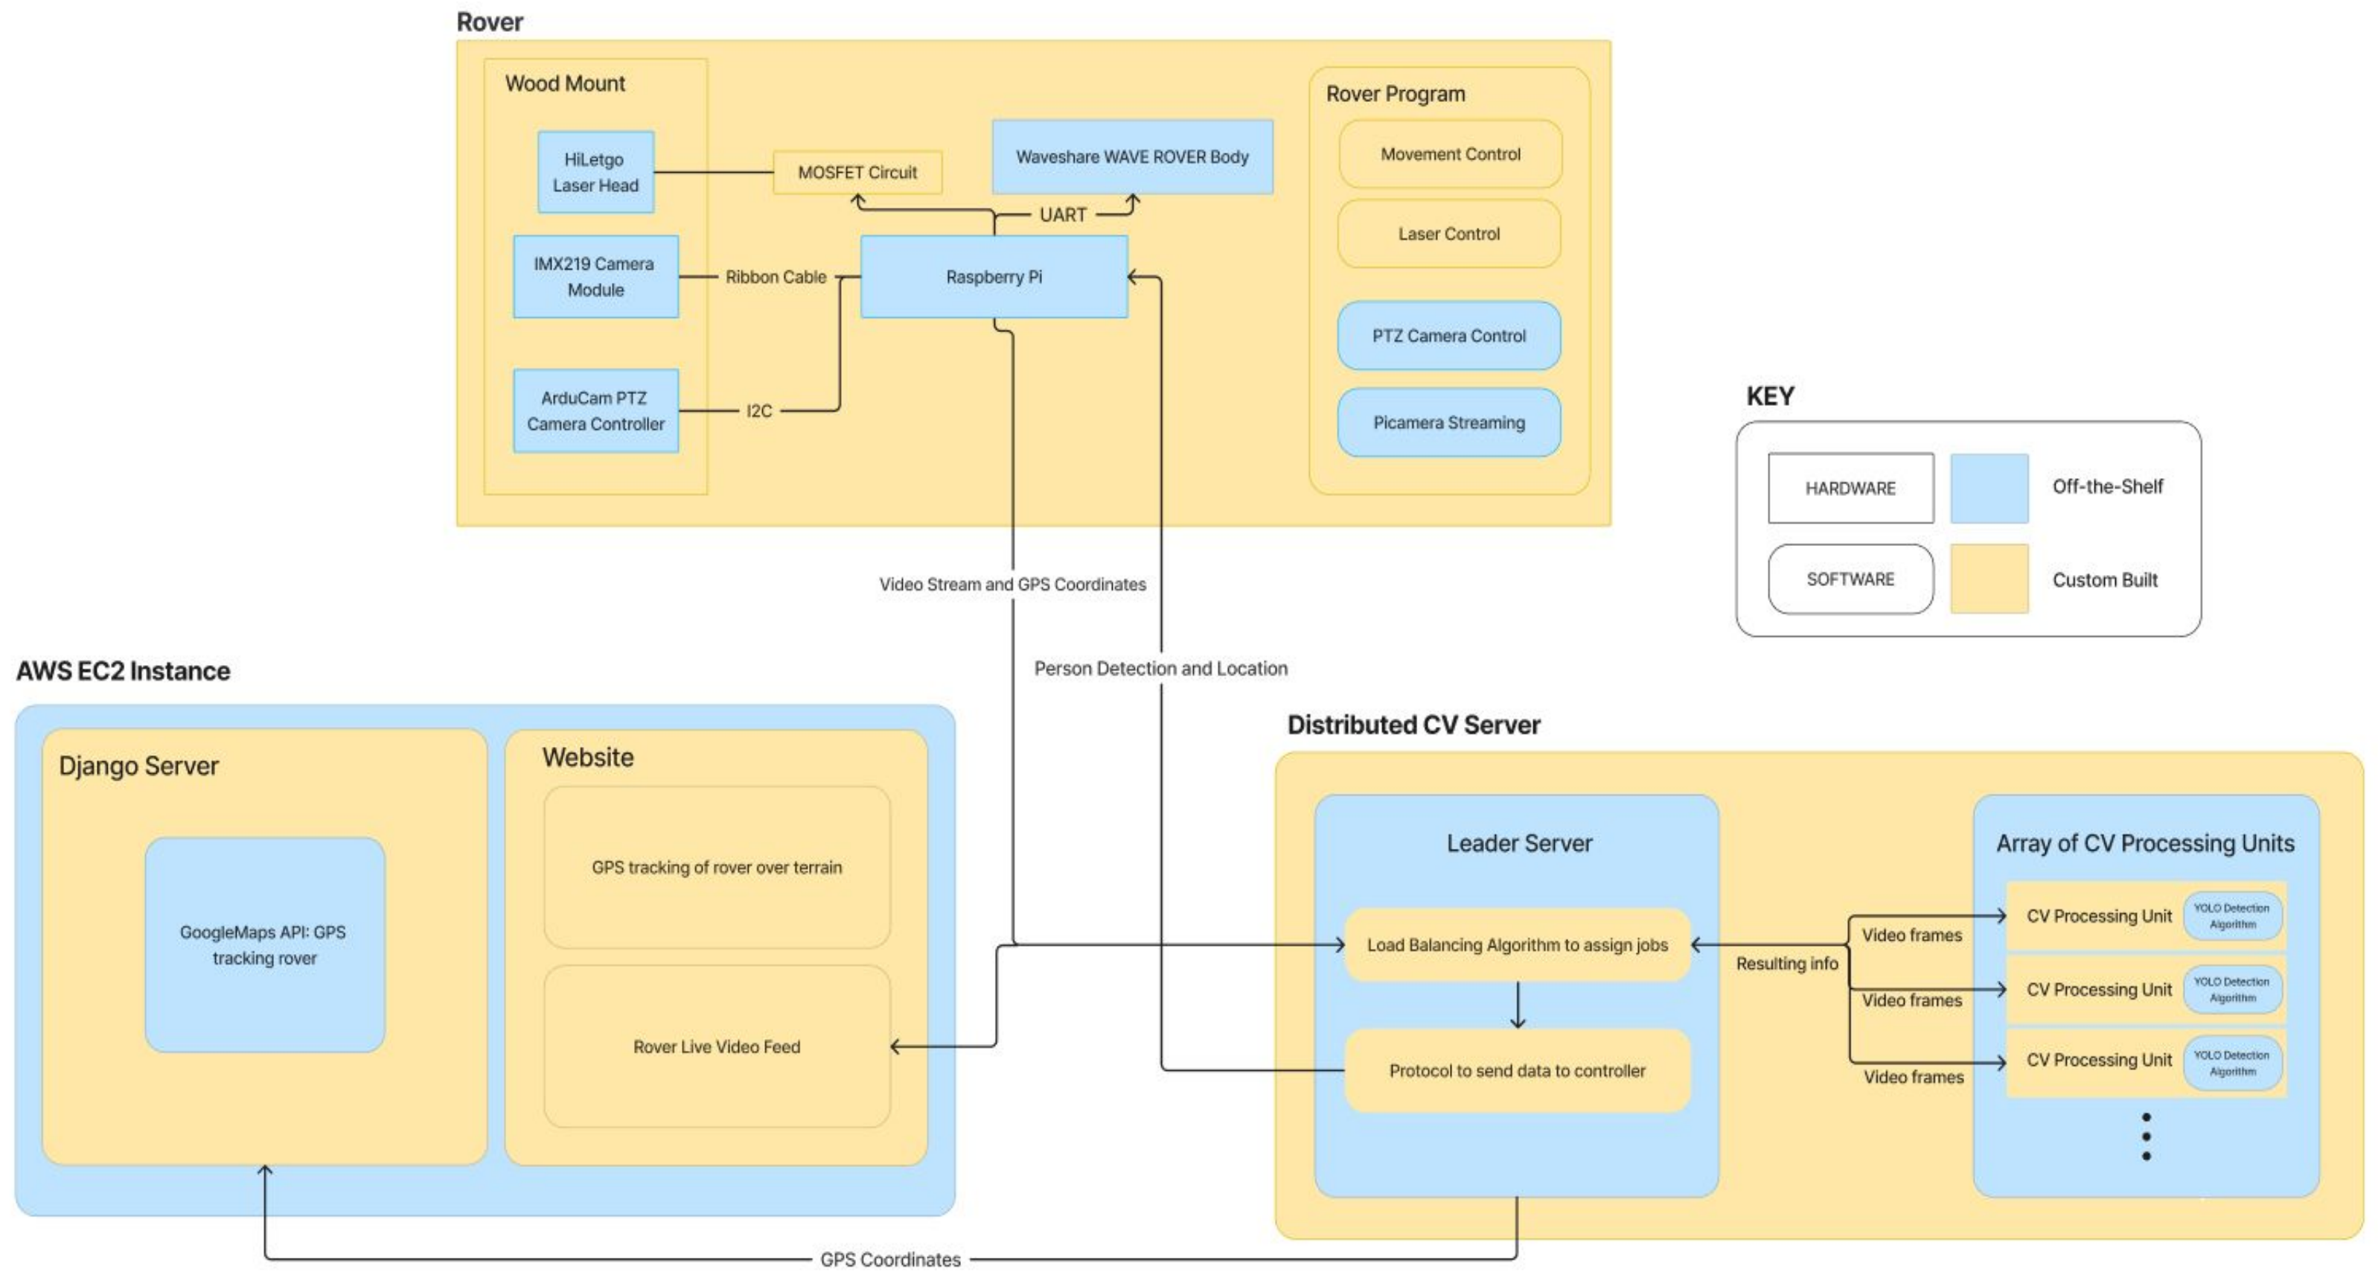

System Architecture Design

Our robot comprises of 5 main subsystems:

(1) Rover: The rover is controlled by a Raspberry Pi and moves autonomously through UART communication from a custom rover program. This allows it to search for people in a predefined creeping line search pattern, and upon being notified that a person was found, switches to a targeted aiming movement.

(2) Camera: The Pan-Tilt-Zoom IMX219 Camera is mounted on top of the rover and displays video data through a ribbon cable on the Raspberry Pi. To increase peripheral vision of the camera, the Raspberry Pi uses I2C to pan the camera 180 degrees horizontally with an Arducam Pan Tilt Zoom controller as the rover is moving. With the Picamera2 library, streaming is enabled and sends live video footage to a MJPEG server that is hosted by the web application and processed by the CV server.

(3) Laser: A laser is statically mounted above the camera, and is powered using the Raspberry Pi. It is positionally directed using the controls for the camera, and is only activated when the camera has centralized the target in its FOV.

(4) Distributed Computer Vision Server: The Distributed CV server runs on the Gates-Hillman Cluster machines and uses synchronization in Golang to accelerate computer vision on video data. A "leader" server node uses Round-Robin load balancing to allocate sampled video frames to 4 "worker" nodes that run YOLOv5 object detection. Results of CV are sent to the Raspberry Pi over WiFi by a "communication" server node.

(5) Monitoring Website: To allow for remote monitoring of the search and shine rover, our website allows for live streaming of the rover's current perspective and real time GPS tracking of where the rover is through Google Maps API.

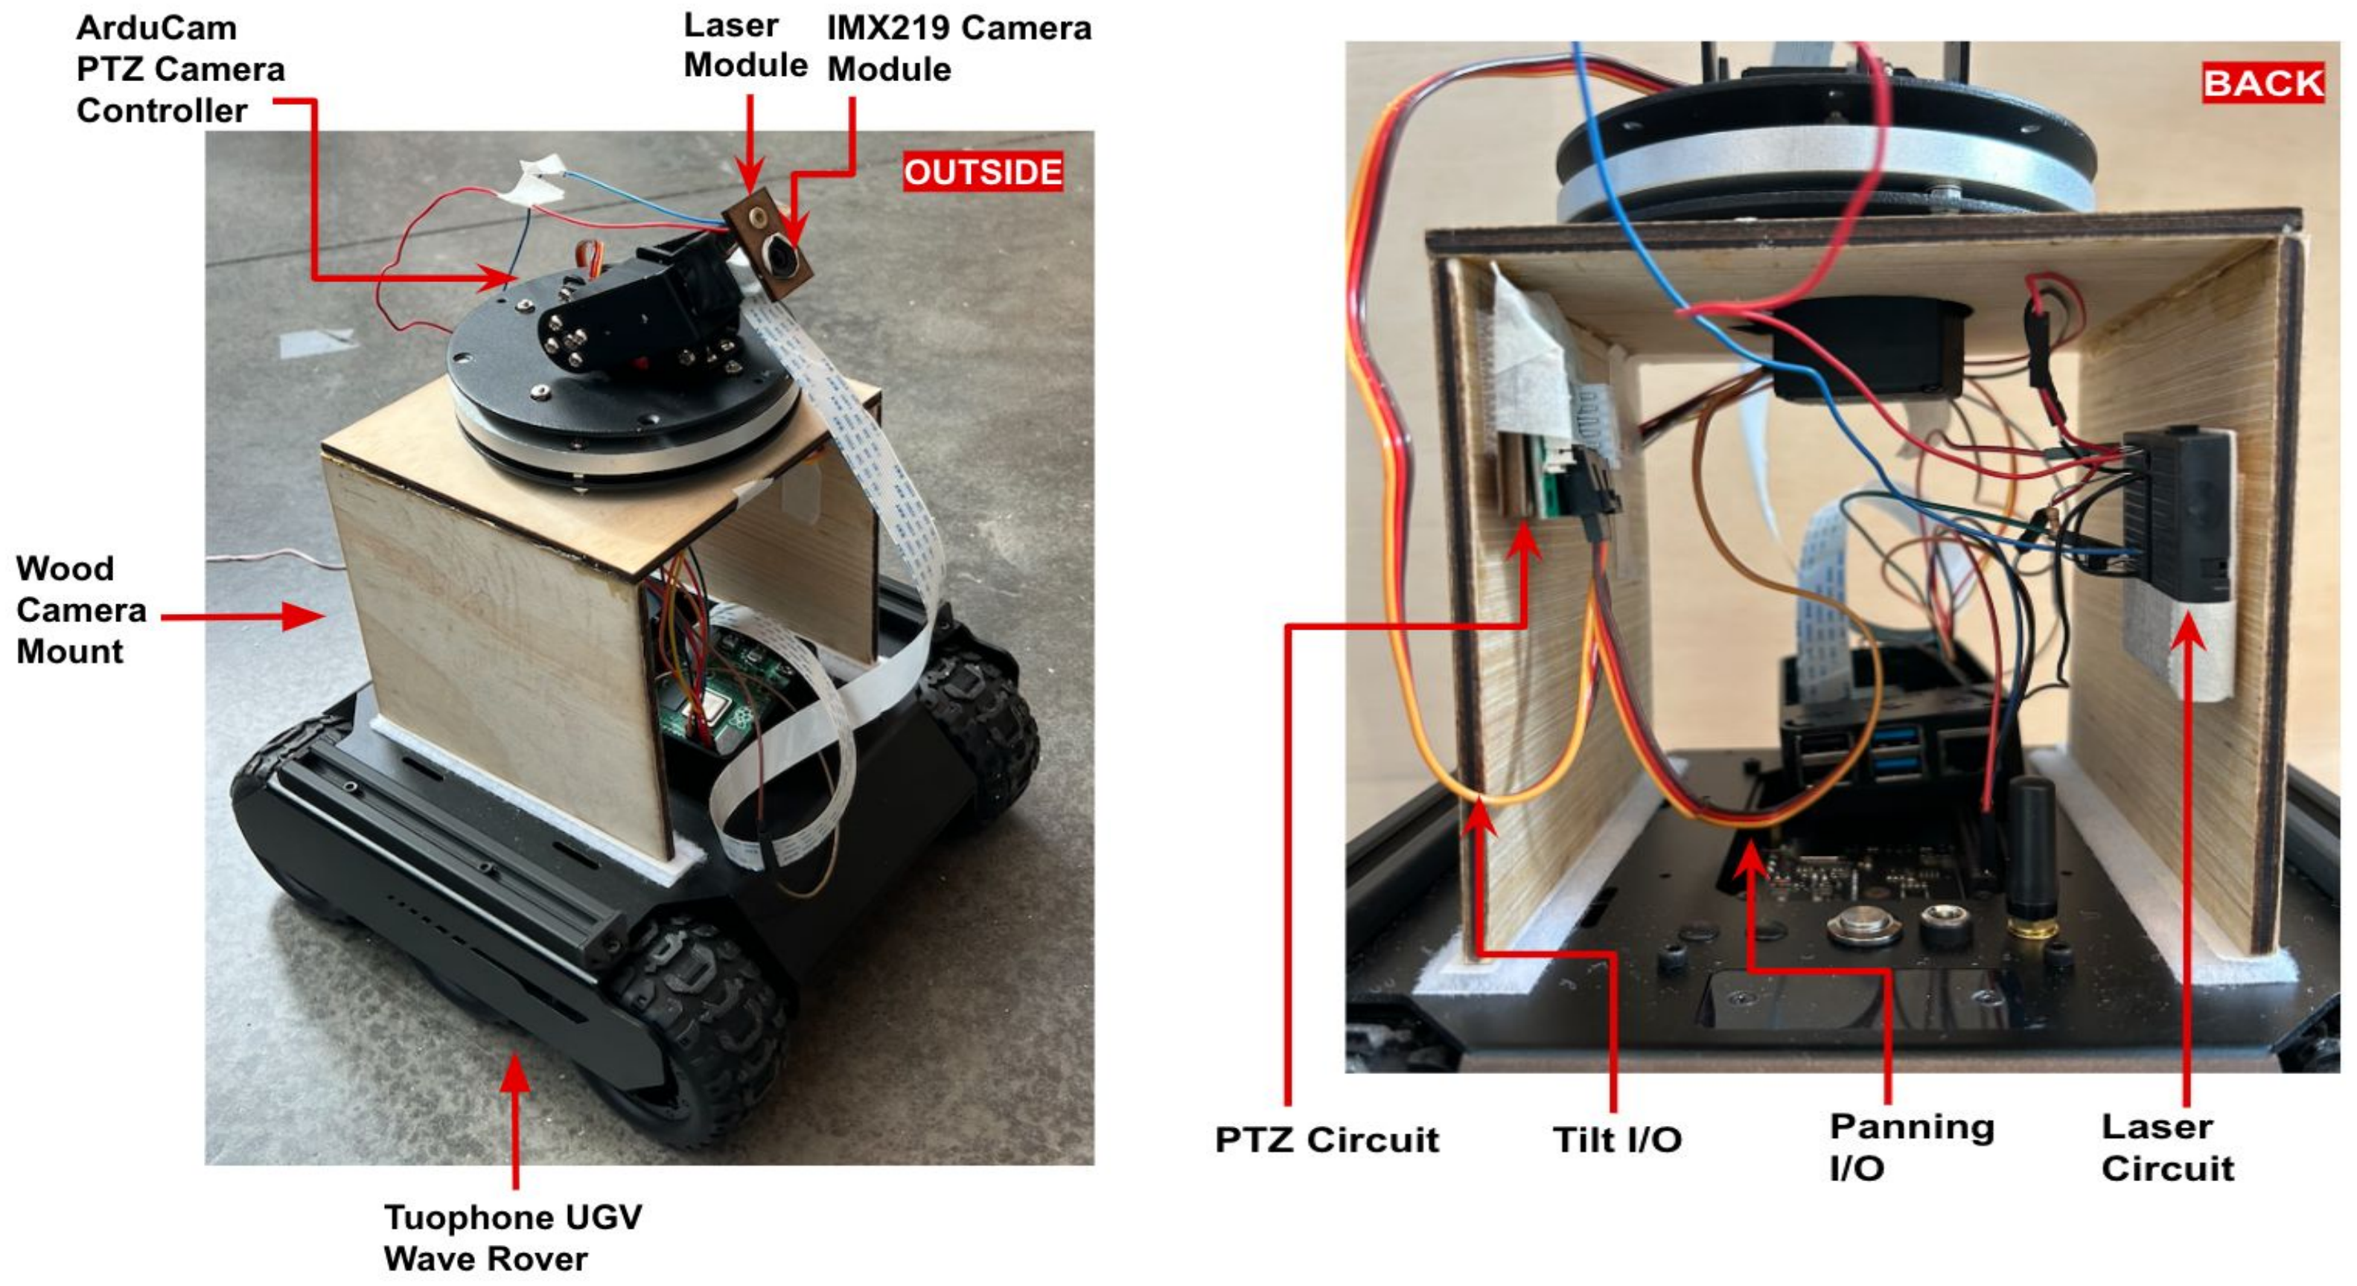

Search and Shine rover: Top and Side Views

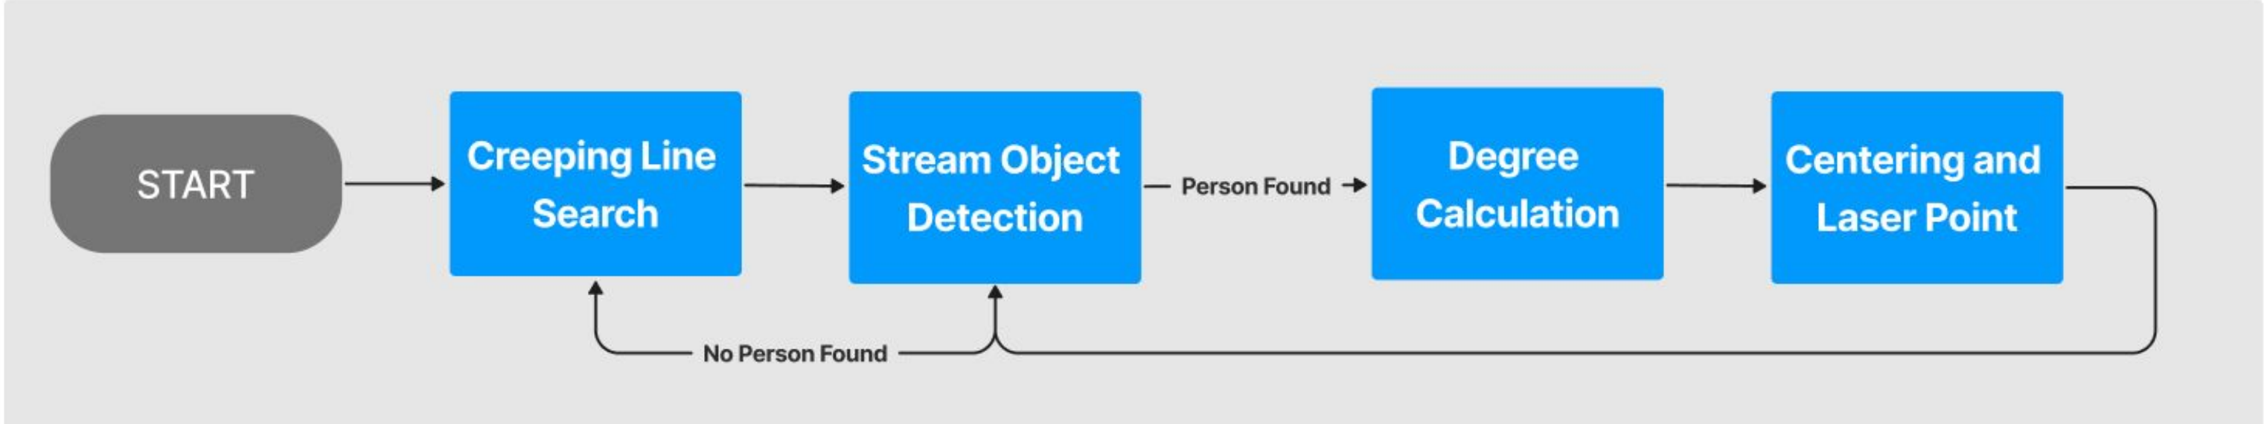

State Machine Diagram

Tradeoffs and Results

A significant design choice was the degree of parallelism to implement in the Distributed CV server. We decided to implement 4 worker nodes after performing speedup analysis. Implementing more nodes led to decrease in speedup due to higher synchronization and overhead costs. Implementing fewer nodes led to insufficient parallelism.

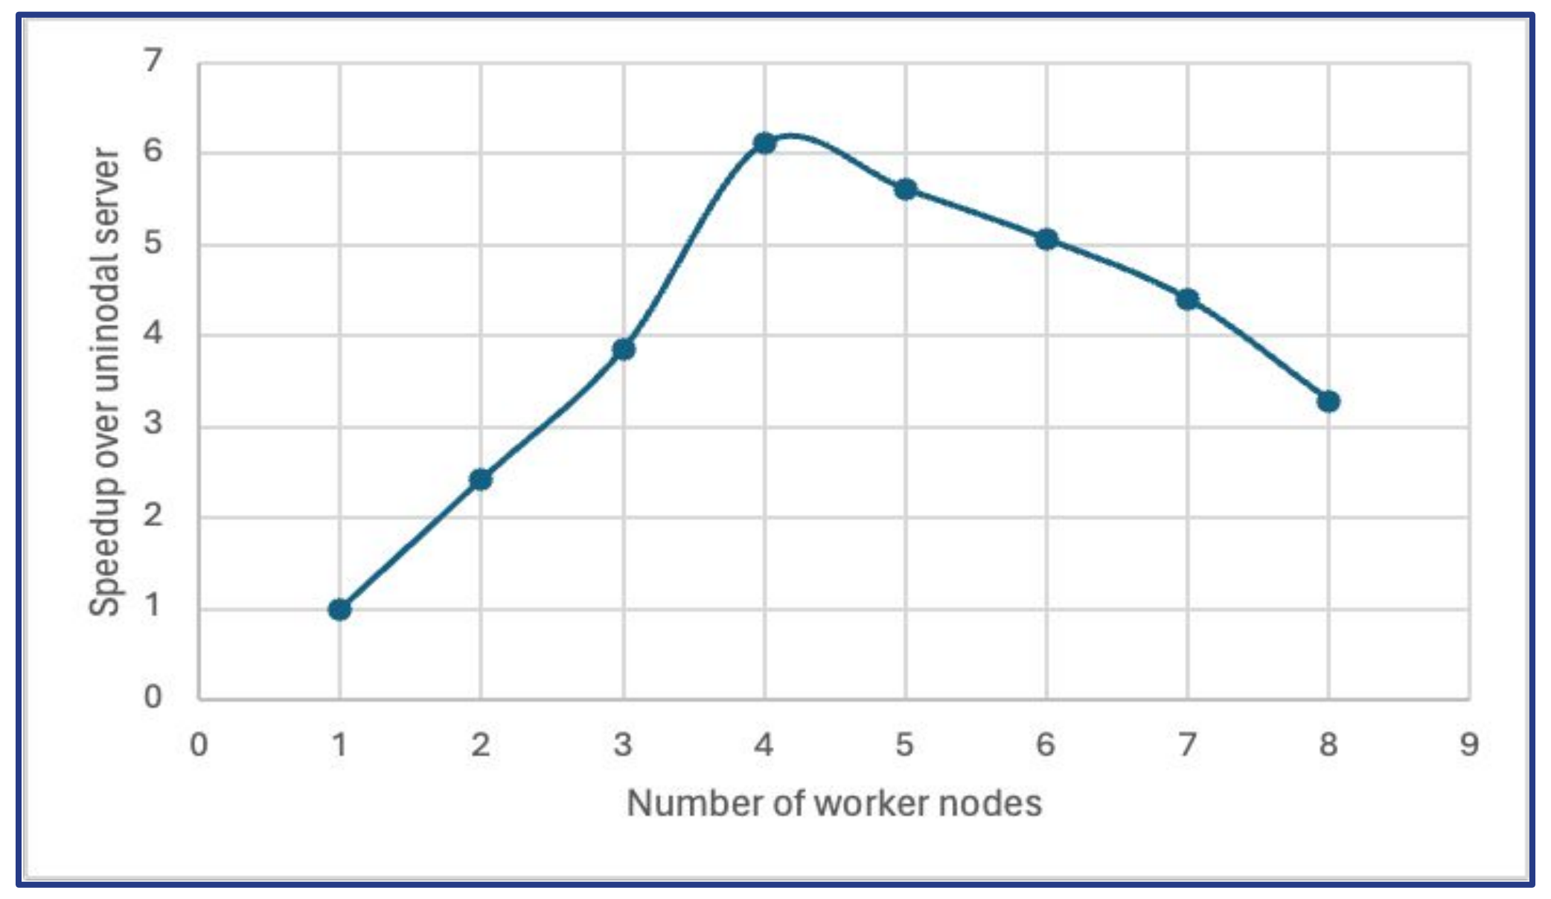

Speedup with increases in server nodes

We performed testing of our system through several unit tests and multiple runs of our scenario tests. The table shown summarizes our test results across 20 sample runs for identifying a human. For these runs, we decided to test 4 different scenarios where no human is in the arena, a human is standing, a human is sitting, and a human is sitting with their face obscured.

As seen below, we achieved most of our use case requirements. Our autonomous rover control, however, was greatly affected due to inconsistencies in the rover's motor control when running the same movement commands. Our power consumption also increased due to high power requirements of the camera, laser, and microcontroller unit; however, we achieved 60 minutes of drive time, which exceeded our requirements. Our system is able to accurately detect humans in different positions and point the laser accurately.

Requirements	Goal Metrics	Results
Accurately identify humans in a flat landscape	Top-1 Accuracy: > 80% Top-5 Accuracy: > 90%	Top-1 Accuracy: **95%** Top-5 Accuracy: **100%**
Autonomous rover control	Can move in pre-specified pattern, and correctly rotate to laser-pointing position if person found (± 1 feet of staying on course)	Average Distance Off-Course: **8.2 feet**
Low latency	Latency of detection, data routing, and result processing: < 5s	Average Time: **1.44s-1.64s**
Point light to person's GPS location	Offset in person location and calculated location: ± 0.5 feet Offset in person location and laser-point location: ± 1 feet	Avg Offset in person location and calculated location: **± 0.3 ft** Avg Offset in person location and laser-pointed location: **± 0.38 ft**
Power consumption	Maintain <5 minute loss of drive time when searching	Average Loss: **8 minutes**

Test results after 20 sample runs of detecting humans

Conclusions

We were able to successfully develop an autonomous search system that could comb semi-flat terrains and use computer vision to automatically detect missing people through a remotely deployed rover that is monitored through a web application. Furthermore, our cost-effective system fulfilled its requirements of maintaining low latency communication between subsystems and boasting a long patrol time by using power efficient search methods.

For next steps, we would like to include ultrasonic sensors to help the rover identify its surroundings. From there, the rover would be able to avoid obstacles autonomously and could extend its use case to an all terrain search.

Due to the low cost and scalability of our rover, we can adapt it to work in hazardous conditions by increasing its durability with weather-resistant hardware and using fire detection software to navigate in wildfire zones for missing people, further exemplifying our mission of improving public safety.

Acknowledgements

We would like to thank our advisor Professor Hyong Kim and teaching assistant Aden Fiol for offering guidance and support during the development of this project. We would also like to thank the Techspark staff and Quinn Hagerty for providing us with resources and space to further our project along.

Spring 2024
18-500 ECE Design Experience