The Illuminator

Authors: Gleb Ryabtsev, Claire Kim, Sidharth Srivastava
Affiliation: Electrical and Computer Engineering, Carnegie Mellon University

***Abstract*—Lidars and time-of-flight cameras sensors are widely used in land and air robotics. However, their application in underwater robotics is limited because they use infrared light, which water blocks. We describe an efficient but compact, low cost design utilizes visible light for our custom time-of-flight(ToF) 3D imaging system for use in underwater robotics. Unlike stereo cameras that require texture-rich scenes for feature matching, ToF directly measures distance through light travel time, making it robust to the low-contrast, featureless environments common in turbid underwater conditions. In this report, we describe a visible-light based ToF system with underwater capabilities and design, and report results that indicate that visible light is a useful medium for ToF imaging.**

***Index Terms*—3D imaging, high speed illumination, LED, LED driver, time-of-flight, ToF camera, underwater imaging, visible light sensing**

1 INTRODUCTION

Autonomous underwater vehicles (AUVs) are increasingly relied upon for tasks like marine research, environmental monitoring, and underwater inspection (for infrastructure, ships, etc). A critical challenge for these systems is perceiving and mapping their surroundings in 3D, especially in murky or low-light environments under the sea where traditional stereo-vision approaches fail. The use case for this project is to equip AUVs (specifically the TartanAUV platform at Carnegie Mellon University) with a compact, low cost, 3D imaging system that enables close range navigation and manipulation. By providing accurate scene depth information underwater, the system allows robotic arms or manipulators to interact precisely with objects on the seafloor or in confined environments.

Current underwater depth sensing technologies each present significant drawbacks. Sonars are often the sensors of choice for long-range perception, but suffer from low resolution and multi-path effects. This is especially problematic for close-range perception for manipulation. Furthermore, phased-array 3D imaging sonars are often prohibitively expensive, with entry-level models priced above $30,000. Stereo-vision offers a more economical alternative, but also suffers from limitations. Stereo-matching algorithms rely on visual features in the images to compute a disparity map. Especially in polluted oceans, underwater scenes often have low contrast, poor illumination, and simple geometry, which makes stereo-matching challenging. Time-of-flight cameras overcome these stereo-vision limitations by actively measuring depth through the round trip travel time of emitted light pulses, rather than relying on passive feature correspondence between two camera views. This active sensing approach provides direct depth measurements independent of the scene texture or ambient lighting conditions. By illuminating the scene with a controlled visible-light sources and precisely timing the return signal, ToF cameras can generate accurate depth maps even in the feature poor, low contrast underwater environments where stereo systems fail. In the air, lidars, time-of-flight cameras or structured light cameras[1] often have better performance than traditional stereo cameras. However, commercially available sensors operate in the infrared wavelength range, which is blocked by water. Therefore, they cannot be used on an AUV.

In this project, we design a time-of-flight camera that can be used underwater. We make it possible by using a **visible light source** rather than an infrared one. Our goal is to create an **integrated sensor module** that can be installed on a wide range of underwater vehicles.

2 USE-CASE REQUIREMENTS

2.1 Applications of AUVs

Autonomous underwater vehicles are widely used in inspection and maintenance of critical subsea infrastructure such as oil and gas pipelines, offshore platforms, and undersea communication cables. They enable routine monitoring of mechanical integrity, corrosion, and biofouling, as well as precise localization of faults or damage without relying on human divers. In addition to industrial use, AUVs support environmental and scientific surveys by mapping seafloor topography, inspecting marine habitats, and collecting visual and acoustic data for long-term studies. These missions often require the vehicle to operate close to structures or the seabed, where accurate, high-resolution 3D perception is essential for safe navigation and manipulation. This motivates the integration of compact time-of-flight (ToF) imaging systems to provide dense, reliable depth information in short-range underwater operations.

2.2 TartanAUV

Given the broad range of use-cases of underwater robots, it is impossible to produce a single set of design requirements that would fit every application. For example, inspecting an large dam for signs of structural failure would require a very different camera than inspecting a ship's propeller. In the first case, one would need wide field-of-view

[1] Eg., *Intel RealSense* products.

and long range to cover a large area quickly. In the second case, a high-resolution close-range sensor is preferable to detect small fractures, dents, corrosion, or biological build-up. Moreover, defining form-factor or electrical interface requirements is difficult because most commercial AUV designs are proprietary.

'Instead of trying to tailor our sensor to a specific *commercial* application, we take a different route. We are collaborating with and taking inspiration from TartanAUV, an undergraduate student organization building underwater robots to compete in the international RoboSub competition. The competition involves a variety of navigation and manipulation tasks that must be completed by an AUV fully autonomously. RoboSub is sponsored by various companies engaged in underwater robotics, as well as the U.S. Office of Naval Research. The competition tasks are designed to be representative of typical problems solved in the industry. Thus, a sensor that can be used effectively in the RoboSub competition scenario in a clear water pool will likely be suitable for numerous commercial applications.

There are two tasks that stand to benefit the most from a close-range depth sensor.

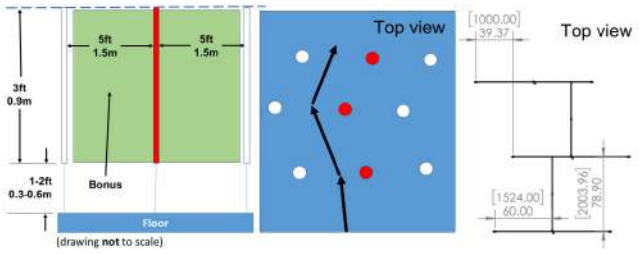

Figure 1: RoboSub 2025 navigation task (view from above). Image courtesy of RoboNation.

The **navigation task** requires the robot to pass through a narrow channel of poles without touching them. The gap between the poles is 1.5 m. (See Figure 1.) To accomplish this task, a 3D image sensor must be able to localize the poles quickly to compute and track an optimal trajectory.

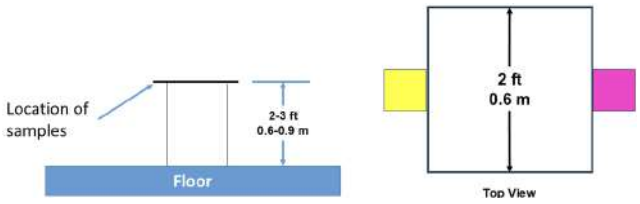

Figure 2: RoboSub 2025 manipulation task (side view). Image courtesy of RoboNation.

The **manipulation task** involves picking up objects resembling plastic waste items from a 2 ft-by-2 ft table, and placing them in bins. This task requires the 3D camera to localize the objects initially, as well as to track the objects as the vehicle and/or the manipulator is guided towards them. The objects are bottle-shaped, around 2 inches wide and 7 inches tall.

Note that the competition is held in a standard Olympic swimming pool with clear fresh water.

2.3 Use-case requirements

The use-case requirements are dictated by the RoboSub competition tasks above, as well as by the physical characteristics of TartanAUV's vehicle, *Osprey* (Figure 3). Osprey moves in the water using eight thrusters and reaches a maximum speed of around 1 m/s. To pick up objects for the manipulation task, Osprey is equipped with a suction tube that is shown in Figure 4. The inner diameter of the suction tube is 4 inches.

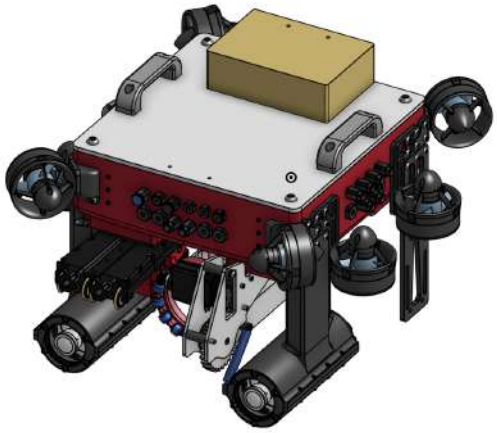

Figure 3: Osprey AUV overview. Image courtesy of TartanAUV.

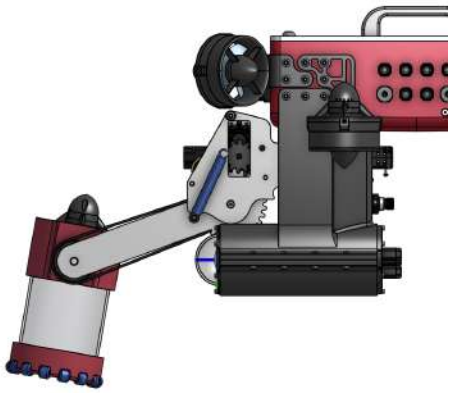

Figure 4: Osprey AUV with the suction tube deployed. Image courtesy of TartanAUV.

We list the use-case requirements below.

- **Maximum range:** 2 **m.** This is the maximum distance at which the sensor should be able to detect an object. This requirement is derived from the navigation task (see Figure 1). When passing between a pair of obstacle poles, the robot must be able to see the next pair of poles to determine its trajectory. The poles are 2 meters away.

- **Minimum range:** 25 **cm.** This is the minimum distance that the sensor can measure. This is approximate distance from the front of the vehicle (where the ToF is mounted) to the suction tube inlet, where an object will be picked up during the manipulation task.

- **Range accuracy:** ±2.5 **cm.** This requirement specifies how accurately the ToF camera can measure distances to objects. The requirement originates from the manipulation task. The objets that must be picked up for the task are 5 cm wide, while Osprey's suction tube inlet is 10 cm wide. If we center the object in the tube, our position estimate must be within 2.5 cm of the ground truth to ensure that the object can be grabbed.

- **System frame rate:** 10 **frames per second.** An overly slow camera would limit the speed with which tasks can be performed. The rate of 10 FPS matches the speed[2] of the onboard imaging cameras, ensuring that the depth images are synchronized with RGB images.

3 Principle of Operation

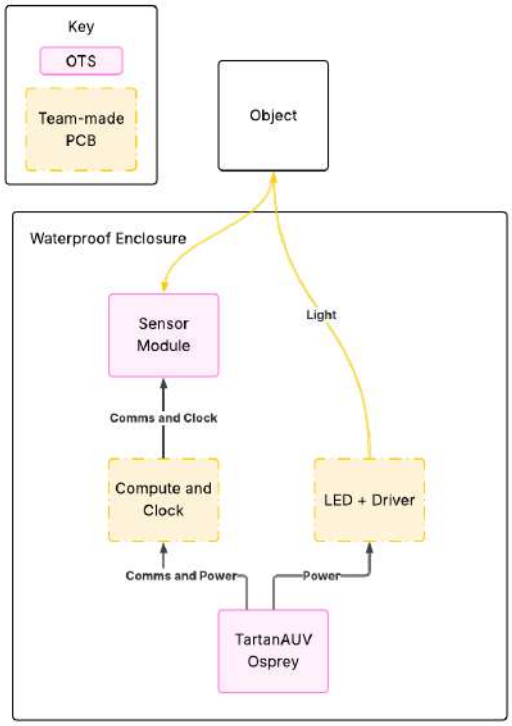

Figure 5: High Level Block Diagram

At a high level, our system consists of two main components: an illumination module and a sense module PCB. The excitation module consists of a high power LED and peripheral driver circuitry to illuminate the scene in front of it. The sense module is an off-the-shelf image sensor which interfaces through a connector to our compute.

The illumination module requires high frequency, high precision modulation of the LED output to the scene. The modulation frequency is calculated based on the desired distance to be measured. During the "off" stage of the LED, the sensor module takes its pixel measurement. The precise calculation for this frequency is shown later in the Design Requirements section.

The sensing module consists of an image sensor with an array of pixels that can read back the intensity of the reflected light from the illumination module. These image sensors can be difficult to integrate with PCBs that we assemble ourselves, so we attempted to look for OTS solutions/carrier boards with peripheral circuitry that are easy to integrate into custom boards. Since members of our team had experience with the EPC660 image sensor/carrier board and it had been used in similar research projects, we decided to build our system around this particular unit.

There were a couple tradeoffs regarding the split of the modules: since the modulation signals might have required communication or other components, we wanted to keep

[2]This rate is limited by the throughput of CV and ML inference pipelines running on the main onboard computer.

this circuitry as close to our compute unit as possible. In addition, we chose to separate out the sensing and illumination modules to ensure that we could vary the distance between the two functionalities to change the overlap of the illumination cone and the sensor FOV. Finally, due to thermal constraints, we had to split the illumination PCB into two different boards, which will be covered in the System Implementation section.

4 DESIGN REQUIREMENTS

Our use-case requirement translate into design requirements on different components of the system. In this section, we will derive these design requirements.

Many of our requirements drive component choices, and those choices affect other design requirements. Therefore, we started by choosing the components with smaller number of alternatives, most importantly the ToF sensor.

4.1 ToF Sensor IC Requirements

The ToF sensor is in many ways similar to a regular camera sensor, and shares a lot of characteristics with it.

Resolution The sensor resolution design requirement is derived from our required lateral resolution and the sensor FOV. With an HFOV of 90, the imaging plane at 40 cm is 80 cm wide. To achieve a 2.5 cm resolution at this distance, we need at least

$$\frac{80}{2.5} = 32$$

horizontal pixels. Vertical resolution is less important because of how the suction tube operates.

Frame rate The sensor must support the required 10 FPS frame rate.

Modulation frequency and phase resolution A continuous wave ToF camera works by modulating the illumination light at a high frequency, and measuring the phase shift of the returned light. The phase is measured by the sensor IC for each pixel. The upper limit on the modulation frequency comes from the maximum range constraint, as the roundtrip time of the light must be less than one modulation period to avoid spatial aliasing. The lower limit comes from the range resolution. A given difference in range corresponds to smaller difference in phase at a lower frequency.

For the maximum frequency, we have

$$\begin{aligned} f_{max} &= \frac{c}{k_w \cdot d_{\max}} = \frac{3 \cdot 10^8 m/s}{1.33 \cdot 2m} \\ &\approx 110 \text{ MHz}, \end{aligned}$$

where k_w is the index of refraction of water and c is the speed of light. For the minimum frequency, we assume conservatively that the sensor has 8-bit phase resolution. Then, we want

$$\frac{2\pi}{2^8} < \frac{4\pi f k_w \Delta d}{c},$$

where Δd is the range resolution. We get

$$\begin{aligned} f_{min} &= \frac{c}{2 \cdot 2^8 \cdot k_w \Delta d} \\ &\approx 17 \text{ MHz}. \end{aligned}$$

The above formulas can be easily derived by considering the spatial wavelength of the modulation wave.

4.2 Optics requirements

Image circle EPC660 pixel field has a diagonal of 8mm. We need a lens with an image circle around 8mm.

HFOV Our lens should provide 90° horizontal field of view.

Depth of field The lens should be able to deliver sharp images of objects at 25cm to 2m range.

4.3 Illumination requirements

Brightness It is intractable to calculate the required illumination brightness theoretically. Any estimate will likely be significantly off due to dissipation effects, unknown reflectivities of target materials, unknown LED attenuation at high frequencies, thermal derating of illumination components, and numerous other factors.

Therefore, we approach this requirement empirically. We know from previous experimentation at TartanAUV that a combined 2000 lumen light source is sufficient to illuminate a scene within 1-2 m, which matches our range requirements. (This was tested in a pool with indoor lighting.)

To account for possible adverse effects of high-frequency modulation, we add a safety factor of 3, and target 6000 lumen as our peak brightness requirement.

Frequency The light source must be able to be modulated at 24 MHz.

5 DESIGN TRADE STUDIES

A few major subsystems of our solution required significant thought to decide between potential options. Chief among these choices include our compute, image sensor, light driver, and modulation clock source.

5.1 Compute

Table 1: Compute Design Matrix

Compute	RAM	Clock (MHz)	Integration	Comms.
H563	640 KB	**250**	**Easy**	**Yes**
H725	640 KB	**550**	Medium	**Yes**
F413	320 KB	96	**Easy**	**Yes**
Jetson	**4 GB**	≥ 1000	Difficult	No

We considered a mix of microcontrollers and SBCs for our compute unit. Our criteria included support for parallel communication with an image sensor (as most image sensors use a similar parallel protocol), RAM size for frame buffer storage, output clock capabilities for modulation, and board footprint/ease of integration. The STM32F4, F7, H7, and H5 series were considered for MCUs, and the Nvidia Jetson Nano was considered for SBCs. Ultimately, we chose the STM32H563 as it had the highest clock output modulation combined with the smallest footprint. The SBC was ruled out as it did not have parallel communication support and was difficult to integrate into a single PCB module, as well as not having a designated timer output.

5.2 Light Driver

Table 2: Light Driver Design Matrix

Light Driver	Circuitry	Integration	Modulation
EPC21603	**Low**	Difficult	**96 MHz**
TPS2816	High	**Easy**	40 MHz

For our light driver, we considered two main options: a high speed MOSFET driver like the TPS2816 or a dedicated LED driver chip like the EPC21603. We considered a few different factors: peripheral circuitry, modulation speed, and ease of integration onto the PCB. The EPC21603 is more difficult to integrate due to its BGA solder package, but it can handle modulation frequencies of up to 100 MHz, while the TPS2816 can only handle up to 40 MHz. In addition, the EPC21603 takes in a differential clock signal directly. As discussed later in the System Implementation section, a differential clock is crucial for signal integrity when travelling between boards. The EPC21603 also does not require any external MOSFETS, which relieved some part selection burden on us. Ultimately, we ended up going with the EPC21603 to fully meet our technical requirements.

5.3 Clock Source

Table 3: Clock Source

Clock	Circuitry	Precision (ps)	Adjustability
Si5338Q	High	**10**	**High**
Internal MCO	**Low**	**50**	Low

We had two main options when considering the modulation clock source: either generating it off of the STM32, or trying to find an off-the-shelf chip to generate the clock. We did a lot of experimentation with both the STM32 TIM and MCO functionalities to output clock signals on the GPIO pins of the MCU. However, we found that at the speeds we needed (up to 96 MHz for the EPC660), the clock division capabilities of the STM32 are severely limited due to their reliance on ARR register values, which made it difficult to precisely control phase offsets on the two different clocks if necessary on the STM32. Thus, we ended up overwhelmingly deciding on an external clock chip, and we were able to find the Si5338Q, which allows for very fast, 96 MHz speeds and picosecond-precision phase offsets, more than good enough for our application.

6 System Implementation and architecture

See the end of this report for a large, system level block diagram.

6.1 Image Sensor Module

The image sensor PCB consists of all of the peripheral hardware necessary to communicate with our chosen image sensor, the EPC660, and transmit the obtained frames to the outside world. The EPC660 has a 320x240 range and can support up to 78 FPS output, which is more than enough to meet our design requirements. The STM32H563 compute unit lives on this board along with all of its decoupling and oscillator circuitry. The STM32 communicates with the image sensor carrier board through the DCMI protocol, which is a parallel, 12-wire data bus with sync lines. All of the image sensor interface lines run through a 60-pin right angle card edge connector that is compatible with the carrier board. In addition, the modulation clock frequencies are generated on this board by the Si5338Q chip. This chip communicates through I2C with the STM32. Two in-phase clock signals are generated from this chip, one for the EPC660 and one for the LED modulation on the other board. The EPC660 clock signal is single ended, as it goes directly to the image sensor through the edge connector. The LED modulation clock signal is differential, as it travels quite a bit farther to the LED driver PCB. This clock runs through a differential pair wire harness for signal integrity purposes to the other LED Driver. Finally, there is an Ethernet PHY chip and another Ethernet connector on this board to communicate with the TartanAUV Osprey submersible. 10V, -10V, 5V, and 3V3 power for the STM32 and EPC660 is derived from the power distribution board.

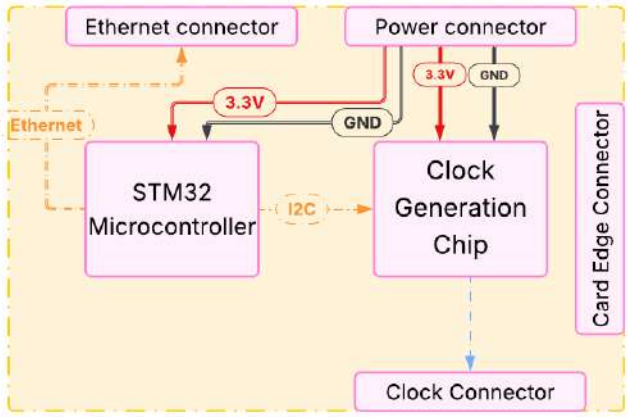

Figure 6: Sense PCB Diagram

6.2 LED Driver Module

The LED driver module takes in the differential clock signal generated from the image sensor PCB. We utilize a high-speed, low-side LED driver chip, the EPC21603, specially designed for LIDAR and ToF applications. We also designed a snubber circuit to mitigate ringing on the gate of the driver and prevent shoot-through problems. [2]After extensive circuit simulations, we determined that the thermal performance of normal FR4 based PBCs was insufficient. Thus, we made the decision to split apart the actual current driver and the LED modulation circuits, with the LED modulation going on an aluminum core PCB that is able to be heat sunk to our overall enclosure.

6.3 Power Distribution Module

The power module will as multiple on-board buck converters, LDOs, and charge pump circuits to generate the -10V, 10V, 5V, and 3V3 signals we need for the the EPC660. It routes all of these power supplies out through a Nanofit connector as well. We decided to create this board during the design phase because we realized that a few of the supplies on the TartanAUV Osprey were not low-noise enough for the sensitive analog signals that the image sensor.

6.4 Firmware

Our firmware implementation for the STM32 is based on interrupts and direct memory access (DMA). On boot, we initialize the EPC660 and Si5338Q chips through I2C commands, starting the modulation clock. This requires over 400 different register writes for proper operation. Then, we will link the DMA controller and set up interrupts for the DCMI peripheral. We also set up the Ethernet peripheral with a UDP client. Most of the actual code logic will be event-driven, as this frees up CPU time and allows for quick responses to communication. The DCMI interrupt will trigger the DMA controller to transfer a frame into SRAM. Using DMA means that the CPU will not be involved with memory transfer at all, allowing it to focus on the Ethernet transmits. Each frame from the EPC660 is 614 kB, which nearly maxes out the entire memory capacity of our compute. We will also have a "half-full" buffer interrupt that will trigger the Ethernet transmits. Thus, we can be transmitting out the top half of the buffer while the DMA controller writes frames to the bottom half of the buffer simultaneously, enabling fast and efficient data transfers [3].

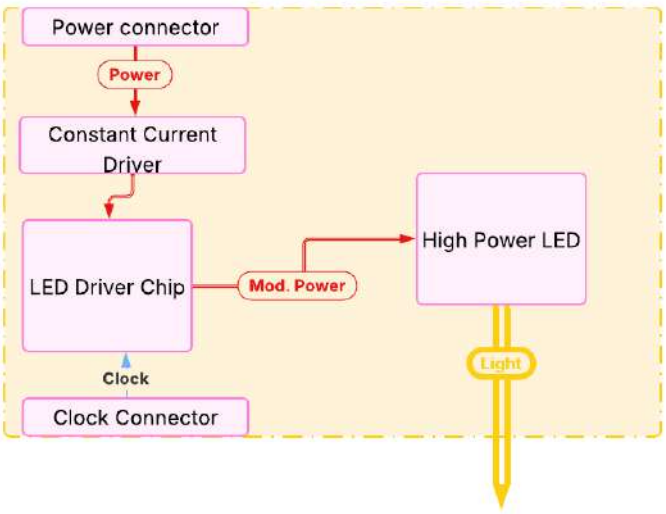

Figure 7: LED Driver Diagram

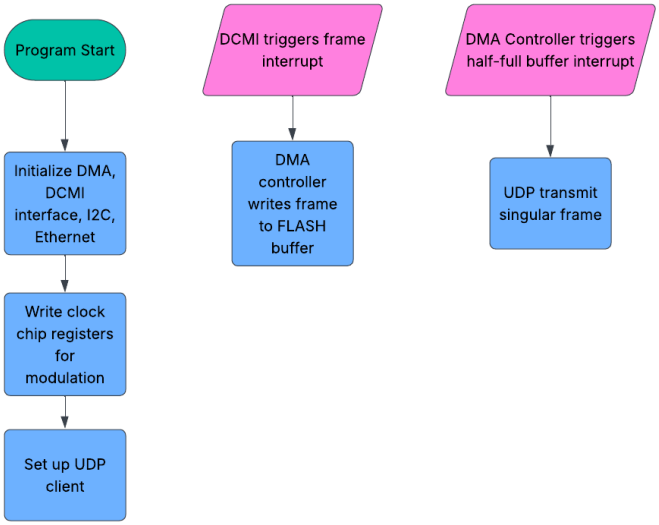

Figure 8: Firmware Flowchart

6.5 Mechanical Integration

Mechanically, we had a few changes by the end from the initial design phase of the project. We had two waterproof enclosures that had been manufactured previously for TartanAUV, complete with O-rings and fastening hardware. We designed our boards to fit within these existing enclosures. In addition, we heatsunk the aluminum core LED modulation PCB to its backing enclosure, which allows us to dissipate the heat from that board directly into the water. Finally, we 3D printed an interface between the lens and the image sensor. Below are CAD screenshots of the final assemblies:

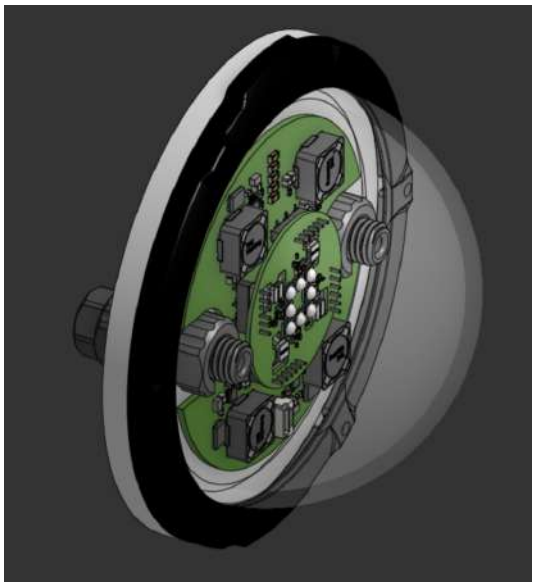

Figure 9: Illumination Module Enclosure

Figure 10: Image Sensing Module Enclosure

7 TEST & VALIDATION

7.1 Results for FPS Test

For this test, we ran our whole data collection pipeline at full tilt for an hour. We used timing functions in Python to determine when a packet was received. We utilized a moving average over 10 seconds to determine the current FPS for a particular time. We found that changing exposure time settings on the EPC660 contributed to variations in exposure time, as shown in the following graph. We were able to successfully tune exposure settings to stay above 10 FPS while getting enough dynamic range between 0 and 2m. Subsequent tests were done at the 25 ms exposure time, and we were able to meet all of the other design requirements with this setting.

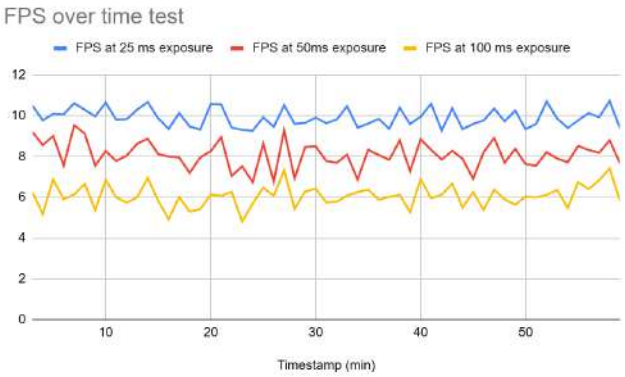

Figure 11: Graph showing FPS over time

7.2 Results for Depth and Error Test

To test depth and accuracy of this sensor, we placed three different cones of the same size at 3 different locations between 0.75 m and 2 m. We then used a laser range finder to determine a precise distance from the sensing element to the cones, and compared that to a normalized depth value calculated from the sensor. We computed a histogram of the pixel values at different depth values to show that different distances are being represented. The average distances shown in the results table here are average distances taken over a square area manually annotated over the cones in a still frame. We also created a depth histogram of all pixels in the frame. The setup and results are shown as follows:

Table 4: Depth Test Results

Laser Distance (cm)	Sensor Distance (cm)	Error (cm)
71.2	73.2	2.0
152.4	154.8	2.4
210.7	213.1	2.4

Figure 12: Cone Test Setup

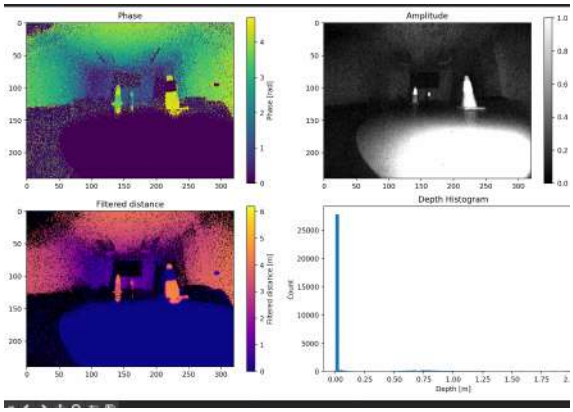

Figure 13: Cone Test Output Images and Histogram

7.3 Results for Resolution Tests

For the resolution test, we took a small hose of diameter 5 cm and showed that it was able to be identified at two different distances of 25 cm and 80 cm. We then repeated the above procedure to average depth values over just the object region in the frame. This test was also completed underwater, which is the reason there are some reflections seen in the overall frame.

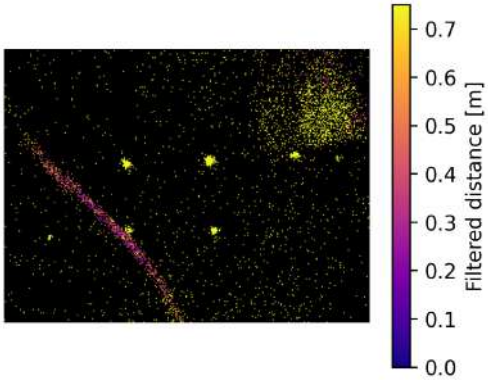

Figure 14: 25 cm Placement of 5 cm diameter object

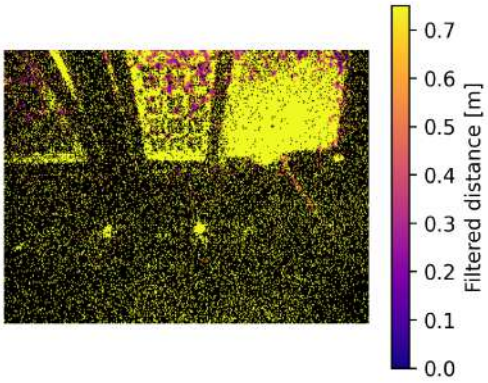

Figure 15: 80 cm Placement of 5 cm diameter object

Table 5: Resolution Test Results

Actual distance (cm)	Sensor Distance (cm)
24.3	25.6
80.3	81.2

7.4 Results for Underwater Tests

We were able to fully test the device under water for over 30 minutes to test full waterproofing, and after inspection there was no leakage. In addition, we tested depth dynamic range under water, and found that pixel density was different for objects at different distances from the image sensor. We employed a similar histogram method to discern depth peaks, and found that the sensor overall also worked underwater.

7.5 Results for FOV Tests

We used a checker board with known width and height of 33.25 cm and 23.5 cm at different depths to calculate a HFOV and VFOV, and we found the HFOV to be 94 degrees and the VFOV to 60 degrees, which tracks nicely with our choice of lens. The field of views were calculated using a distance of 18 cm using the formula $2 * atan2(checkerboarddist/(checkerboarddimension/2))$

8 PROJECT MANAGEMENT

8.1 Schedule

From our design report to now, our schedule has had minor changes. We took an additional week during our design process to go over simulation results and ensure design choices of our PCBs so no additional iterations would need to be made. Then, an unexpected shipment delay of our parts occurred, which we took that week to push forward tasks that could be done sooner than planned. We spent more time in the following week to finish the tasks that we fell behind on, like assembling the boards and unit testing

them with firmware, which helped us stay on track for the rest of the timeline. The schedule is shown in Fig. 17.

8.2 Team Member Responsibilities

We split team member responsibilities broadly based on our previous experience and the large subcomponents of our whole assembly. Gleb worked on the analog driver PCB for the excitation LED due to his previous experience in high frequency switching power design. Claire designed the digital logic and sensor PCB to interface with the EPC660 given her past experience with STM32 boards, and Sid wrote the firmware for DCMI and Ethernet interfacing. Unit testing of each of the subsystems was the responsibility of the person responsible for it, while integration testing and mechanical integration was a very collective responsibility. We all contributed to testing and integrating all the individual components and subsystems, as well as a lot of the debugging.

8.3 Bill of Materials and Budget

See the end of the report for a full BOM.

8.4 TechSpark Usage

We utilized tools from the PCB fabrication located in TechSpark. The xray was used to validate small component placements and good connections from parts to pads of the PCB's. Differential probes were available to use to verify waveform of the LVDS clock modulation signal generated from our clock chip. The reflow oven was also used for PCB assembly.

8.5 Risk Management

To ensure reliable operation and minimize electromagnetic interference (EMI) in our ToF illumination system, several design strategies were made. For our PCBs, high frequency signal integrity was prioritized through single-layer routing with minimal vias, keeping critical clock signal traces as straight and short as possible. Component selection focused on low equivalent series resistance (ESR) capacitors in compact 0603 to 0402 packages to minimize parasitic effects and for ease of manual PCB assembly. Power delivery was carefully managed by placing decoupling capacitors immediately adjacent to their corresponding integrated circuit (IC) pins on the same layer, each with dedicated vias, while crystal oscillators were positioned directly next to their respective pins to minimize clock signal degradation. A continuous solid ground plane on the second layer immediately below the signal layer provides a stable reference and return path for all signals. To maintain signal quality across differential pairs and time-critical paths, impedance matched traces with similar lengths and termination resistors. Finally, bulk tantalum capacitors filter input power supplies for overall system stability, while ferrite beads isolate sensitive analog power pins from digital switching noise, ensuring clean operation of precise timing circuits essential for accurate depth measurements.

As for our schedule, we followed our Gantt chart outlining all major milestones, complete with individual tasks leading up to integration of our subsystems. This structured approach allowed us to work in parallel for our allocated responsibilities as well as manufacturing bringup to allocate resources appropriately to prevent bottlenecks. Recognizing that hardware development is prone to unforeseen challenges such as parts shipment delays and design iterations, we built in a week of flexibility as a buffer for these unexpected issues. Regular progress reviews against the Gantt chart enabled the team to proactively identify slipping tasks and maintain overall project momentum with a plan of action in place, ultimately delivering our system within the final demo deadline despite the minor setbacks mentioned before in 8.1.

Resources management was handled by making use of TartanAUV's inventory of parts that did not need to be purchased, reducing cost in our budget to spend on components that could not be borrowed. Several bodges on our PCBs were made to avoid reiterations of the boards, as tariffs and shipment took up a considerable amount of money. Maintaining strong communication across the team allowed for transparency and schedule adjustments based on personnel availabilities and unexpected delays.

9 ETHICAL ISSUES

Several edge cases present challenges for our ToF underwater imaging system. From an environmental perspective, high-power illumination in sensitive marine ecosystems poses as a concern as our visible-light ToF system requires active illumination to function. And in ecosystems with light sensitive marine life could cause behavioral disruption, interfere with natural bioluminescence communication or damage photosensitive organisms. Additionally, the system's performance may degrade in extreme turbidity conditions, potentially leading operators to increase illumination power beyond safe operation levels. Partnerships with marine conservation organizations can develop certification standards for ecosystem safe operation. Because our system is not designed for deep sea environments as of yet, artificial light pollution will not compromise light sensitive species within undisturbed deep-sea environments. But to minimize these harms, environmental impact assessments across diverse ecosystems before commercial deployment will help establish guidelines.

From a security standpoint, the most concerning edge case involves malicious actors deploying our technology on weaponized or sabotage oriented AUVs. The precise 3D mapping and close range navigation capabilities that we meant for infrastructure inspection could equally enable targeted attacks on critical undersea structures. They can use it for striking these vulnerable targets in international waters where accountability and jurisdiction is unclear. This would threaten millions globally who rely on

undersea internet cables for connectivity. Includes remote workers, international businesses, and entire nations dependent on a small number of cable land points would be devastated if these bottleneck points were damaged. A single severed continental link could isolate entire regions form global connectivity, especially in the modern day where everything seems to rely on internet communications. To prevent this, rigorous know-your-customer protocols must be put into place, work with export control authorities to classify our device under dual-use technology regulations, and implement technical safeguards like mandatory serial number registration. Contracts with AUV manufacturers should ensure that our technology exists only in platforms with audit capabilities, while partnering with international maritime security organizations will enable threat intelligence sharing and rapid response if it is identified that our system is being used for suspicious activities.

10 RELATED WORK

There has been some research related to visible-light ToFs, and is some of where we got our inspiration for this project. There was an IEEE Oceans Conference paper that used this image sensor unit that we are using [1]. However, their design was very clunky and built around existing OTS compute units, so we wanted to build a similar, yet simpler system from the ground up.

11 SUMMARY

To conclude, we were able to successfully create a 3D TOF system that met all 3 of our design and use case requirements on range, resolution, and FPS. We were able to gain insights into the nuances of optoelectronic systems, as well as effective strategies for intricate PCB assembly.

11.1 Future work

In the future, we hope to actually calibrate these cameras on known calibration checkerboards, which are tuned to the optical parameters of a particular camera. This would also involve fixturing the distance between the light modulation as well as the image sensor to keep optical properties constant. We also want to try and We hope to complete this over the next semester, as well as potential discussions about commercialization.

11.2 Lessons Learned

We learned quite a few lessons, including the potential issues with manual PCB assembly. If we respun the boards, we would get them out of house assembled to avoid issues. In addition, we learned about new modes of operation of memory controllers on microcontrollers, and how to push a particular chip to its limit. However, we were able to mitigate a lot of issues by careful design and part selection, which is something we will carry forward to future career projects.

Glossary of Acronyms

- AUV - Autonomous Underwater Vehicle
- DCMI – Digital Camera Messaging Interface
- DMA – Direct Memory Access
- OTS – Off the shelf
- FPS – Frames per second
- HFOV – Horizontal Field of View
- VFOV – Vertical Field of View

References

[1] Evelyn S. DeMember et al. “Underwater Time-of-Flight Camera for Remotely Operated Vehicle Applications”. In: *(MTS/IEEE OCEANS Conference)*. Manuscript, accessed via GCC: `https://www.gcc.edu/Portals/0/MECE-LR_ToF-Camera-IEEE-Oceans-2023-Manuscript.pdf`. 2023.

[2] Larry Li. *Time-of-Flight Camera — An Introduction (Technical White Paper)*. Tech. rep. SLOA190B. Revised May 2014, `https://www.ti.com/lit/wp/sloa190b/sloa190b.pdf`. Texas Instruments, 2014.

[3] STMicroelectronics. *How to Use the GPDMA for STM32 MCUs (Application Note)*. Tech. rep. AN5593. Rev. 7, `https://www.st.com/resource/en/application_note/an5593-how-to-use-the-gpdma-for-stm32-mcus-stmicroelectronics.pdf`. STMicroelectronics, 2025.

Table 6: Bill of Materials

Manufacturer	Item	Quantity	Price Per Unit	Source
ESPROS Photonics	EPC660-CC Carrier Board	1	$91.63	purchased
ESPROS Photonics	EPC21603 LED Driver	4	$2.91	purchased
Texas Instruments	TPS92515 CC Driver	4	$2.17	purchased
Skyworks Solutions	SI5332 Clock Generator	1	$13.47	purchased
Cree	XPE-GR LED	12	$2.30	purchased
Commonlands Optics	CIL-034 Lens	1	$39.00	purchased
Various	Passive Components		$0.00	TartanAUV supply
JLCPCB	PCB manufacturing		$215.00	purchased

Total: **405.22**

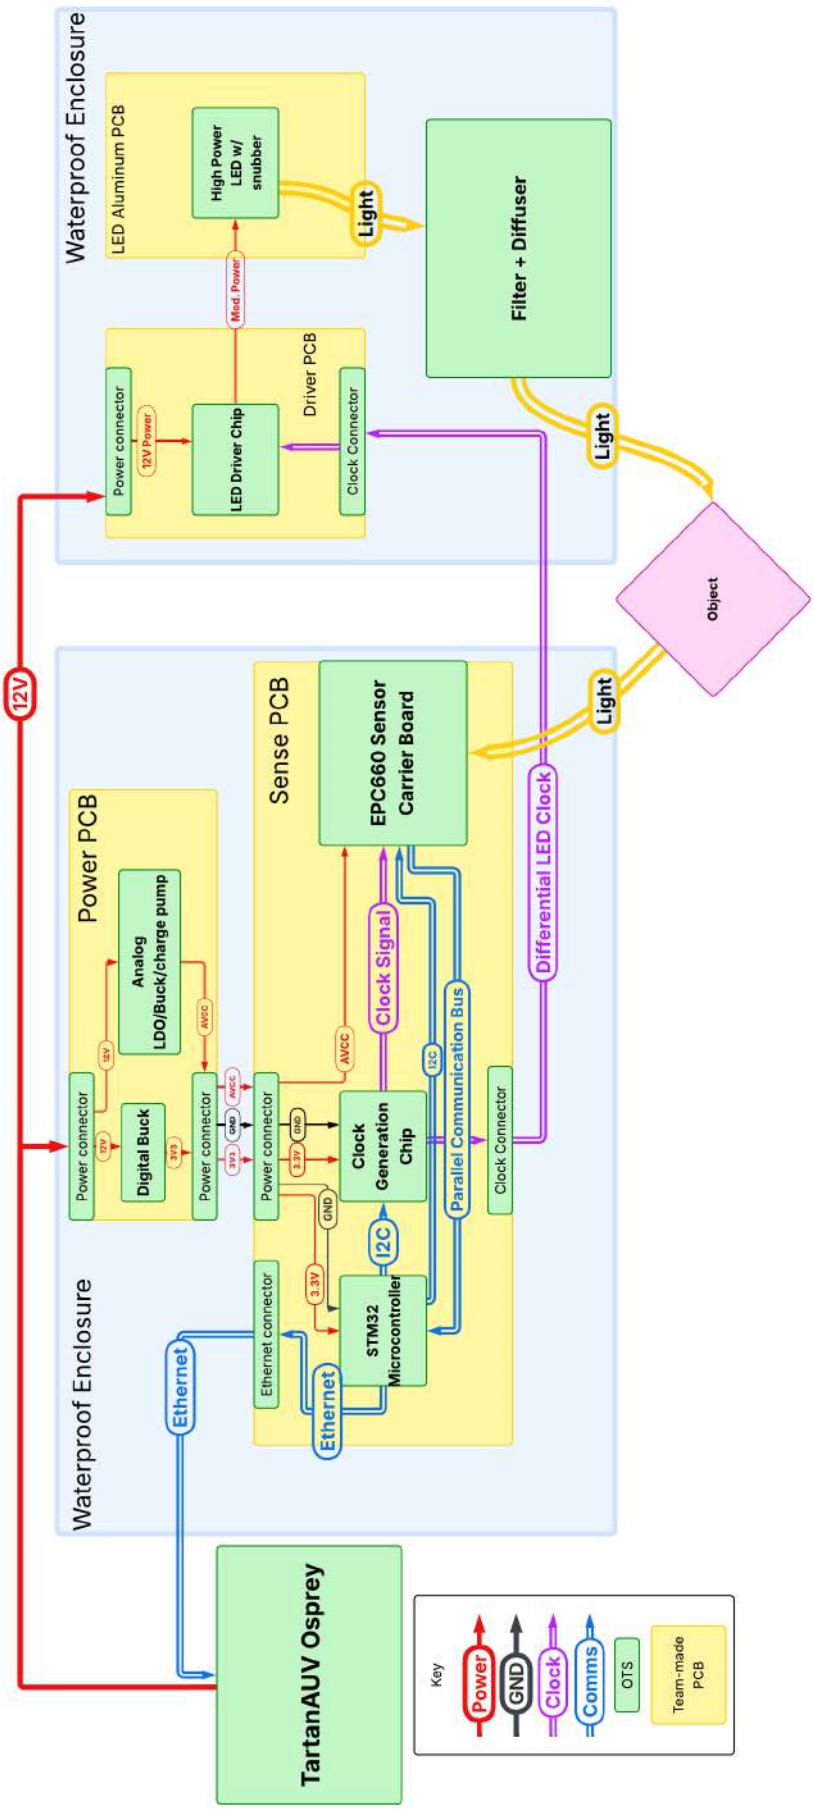

Figure 16: A full-page version of the same system block diagram as depicted earlier.

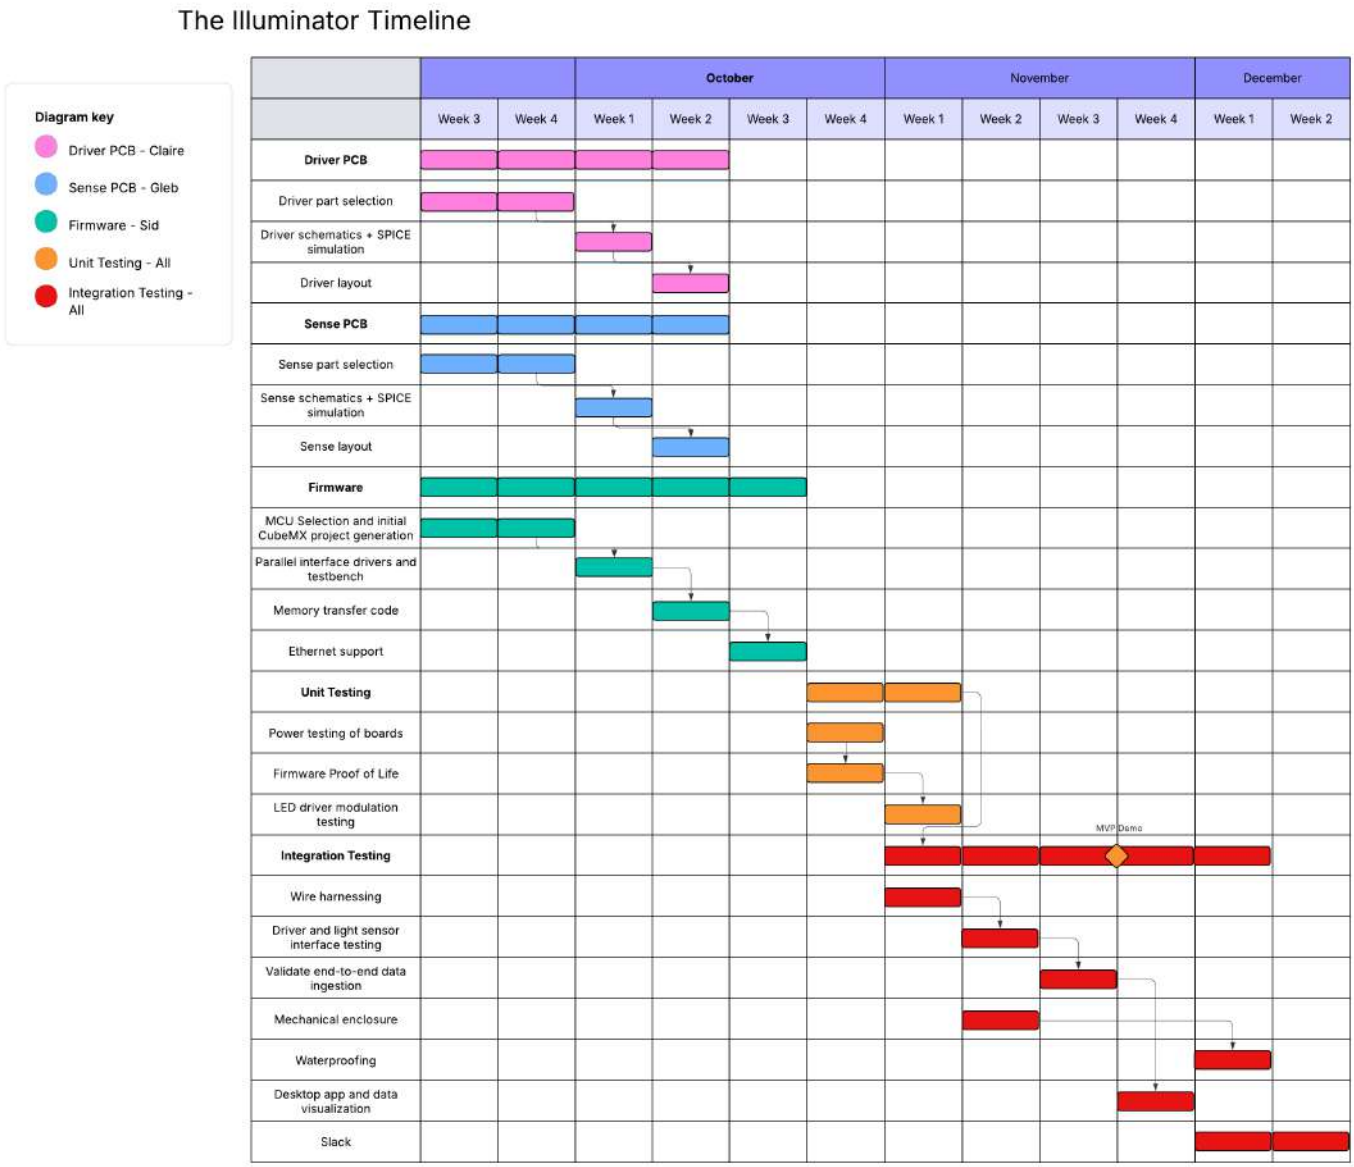

Figure 17: Gantt Chart